APPLICATION NOTE

BRANCH OFFICE CONNECTIVITY GUIDE

Juniper Networks Design Practices for Connecting Branch Offices to Data Centers over a Single Converged Network

Copyright © 2010, Juniper Networks, Inc.

Table of Contents

Introduction . . . 1
Scope . . . 1
Target Audience . . . 1
Design Considerations—Connectivity at the Branch Office . . . 2
Branch Office Reference Architecture—A High-Level Description . . . 2
Branch Office Connectivity over IPsec VPN . . . 3
Design Recommendations . . . 3
Routing Architecture . . . 4
Tiered Network Approach—Layering the Network for Scalability . . . 4
Routing Information Protocol . . . 5
Using RIP with On-Demand Circuits Extensions (Advantages and Disadvantages) . . . 6
Traffic Load Balancing for Type B and Type C Branch Deployments . . . 6
Load-Balancing Solutions in Relationship to Branch Connection Types . . . 7
Using Border Gateway Protocol for Large Networks . . . 7
Using Border Gateway Protocol with Route Reflectors . . . 7
Using Static Routes . . . 8
Using OSPF for Small Number of Branch Offices . . . 9
Additional Design Considerations . . . 9
Using Auto Connect VPN to Create Branch-to-Branch IPsec Tunnels . . . 10
Next Hop Resolution Protocol . . . 10
The Tunnel-Building Process . . . 11
Auto Connect VPN Design Considerations . . . 11
High Availability for the Branch Office . . . 13
High Availability Requirement Levels (Link, Device, Device and Link Levels) . . . 13
High Availability Functionalities . . . 13
High Availability for Branch Office Type A . . . 14
VPN Security Zone Configuration for Type A . . . 14
High Availability for Branch Office Type B . . . 15
Using Secure Services Gateway for Type B . . . 15
High Availabilty for Branch Office Type C . . . 16
Connectivity at the Data Center . . . 18
Implementing a High Availability Enterprise Network at the Data Center . . . 18
Design Requirements . . . 18
Data Center Network Architecture . . . 20
Internet Connectivity . . . 21
Internet Connections . . . 22
Border Gateway Protocol or External Border Gateway Protocol Operation . . . 22
Edge Routers . . . 22
Firewalls—Internet and IPsec . . . 23
Internet Firewalls . . . 24

Copyright © 2010, Juniper Networks, Inc.

VPN Firewalls . . . 27
Shared Services Core . . . 30
High Availability Design of Firewalls . . . 30
Quality of Service Design Requirements . . . 30
WX Series Design Requirements . . . 31
Summary . . . 31
Appendix A Related Documents . . . 31
Appendix B Naming Conventions . . . 32
M Series and J Series Interface Naming Conventions . . . 32
M Series Interface Names . . . 32
Logical Part of an Interface Name (M Series Multiservice Edge Routers) . . . 32
J Series Interface Names . . . 32
SSG Series and ISG Series Interfaces . . . 32
Appendix C Products . . . 33
Product Tables . . . 33
Product Recommendation Based On Branch Office Profile . . . 33
Partner Products . . . 33
Symantec . . . 33
Kaspersky . . . 34
SurfControl and Websense . . . 34
Avaya IG550 . . . 34
About Juniper Networks . . . 35

Copyright © 2010, Juniper Networks, Inc.

Table of Figures

Figure 1: Connecting branch offices, campus locations, and data centers over a single converged network 1
Figure 2: Branch office reference architecture. 2
Figure 3: Multi-tiered/layered network architecture. 4
Figure 4: Two-tier network design for data centers 5
Figure 5: Branch with dual internet connections (load balancing using ECMP). 7
Figure 6: BGP routing design 8
Figure 7: Star topology – connecting branches to central hub 10
Figure 8: AC VPN provisioned tunnels between branches in the same region. 11
Figure 9: Multi-tier topology 12
Figure 10: HA configuration for Type A 14
Figure 11: VPN security zone configuration for Type A 14
Figure 12: Type B optimized – HA configuration. 15
Figure 13: Type B – security zones. 16
Figure 14: Type C – HA configuration. 17
Figure 15: Intra-branch using OSPF 17
Figure 16: Branch Type C – security zones 18
Figure 17: Enterprise network for the data center 21
Figure 18: M Series Multiservice Edge Routers 22
Figure 19: Internet firewalls 25
Figure 20: VPN firewalls 29
Figure 21: VPN firewall IPS policy 29

Copyright © 2010, Juniper Networks, Inc.

Introduction

Designing and scaling an enterprise network for assured network connectivity between branch offices and data centers is a challenge that faces every high-performance organization. This guide can assist organizations to design and implement a secure and reliable enterprise network infrastructure.

Because most enterprises typically employ more users in branch offices than at headquarters, they need a network infrastructure that performs as well as the networks in headquarters and one that delivers secure and assured connectivity. Branch offices usually connect directly to headquarters using either a private WAN link or a VPN over the Internet, or they deploy VPN over the private WAN link. As more branch offices connect directly to the Internet—rather than backhauling Internet traffic to headquarters—this trend introduces a new set of security, performance, connectivity, and reliability challenges.

Scope

This guide provides design guidance for deploying a converged network solution that connects branch office locations to the data center, as shown in Figure 1. It offers the connectivity practices and guidelines for network routing, implementing high availability (HA) options for assured branch office performance and a related HA data center design. All of these topics support a common design goal of connecting multiple branch offices at a maximum of 1,000 locations using an IPsec VPN overlay.

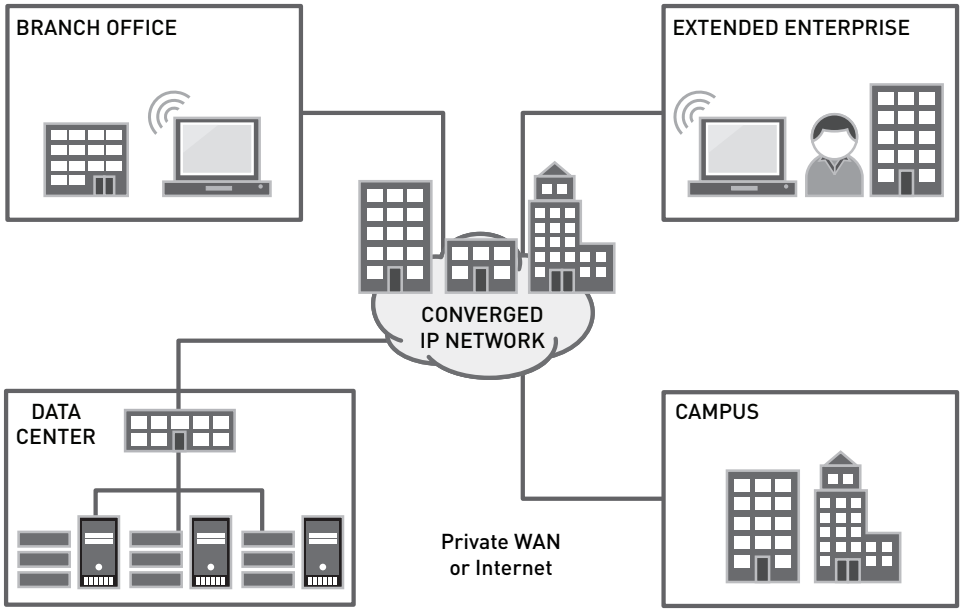

Figure 1: Connecting branch offices, campus locations, and data centers over a single converged network

For each topic discussed in this guide, a corresponding application note is offered that contains additional "how-to" information and device configuration steps for implementing that aspect of the design.

In addition to this guide, three additional guides address enterprise network security, operations, and performance—all leveraging the same connectivity design model. Appendix A provides a list of these guides in addition to application notes, white papers, design documents, and other related information.

Target Audience

- IT managers
- Systems engineers
- Network analysts and engineers
- Network administrators
- Security managers

Copyright © 2010, Juniper Networks, Inc.

Design Considerations—Connectivity at the Branch Office

In this section, design guidance for the following major topics is presented:

- Implementing branch office connectivity using an IPsec VPN overlay
- Using RIP as the preferred routing protocol for the solution
- Employing address traffic load balancing
- Considering additional routing protocols other than RIP for the same design model

After defining the basic design for branch office connectivity, guidance and solutions are presented for configuring HA and fault tolerance for all three types of branch office profiles: Type A, Type B, and Type C. The first section defines the enterprise network architecture at a high level.

Branch Office Reference Architecture—A High-Level Description

The Juniper Networks® Branch Office Reference Architecture categorizes branch office architecture into three different branch office profiles (see Figure 2). Table 1 summarizes the features the branch office profiles and the services they provide. This architecture is used as the basis for the discussions and design practice recommendations. For details concerning the branch office reference architecture, refer to the Branch Office Reference Architecture paper.

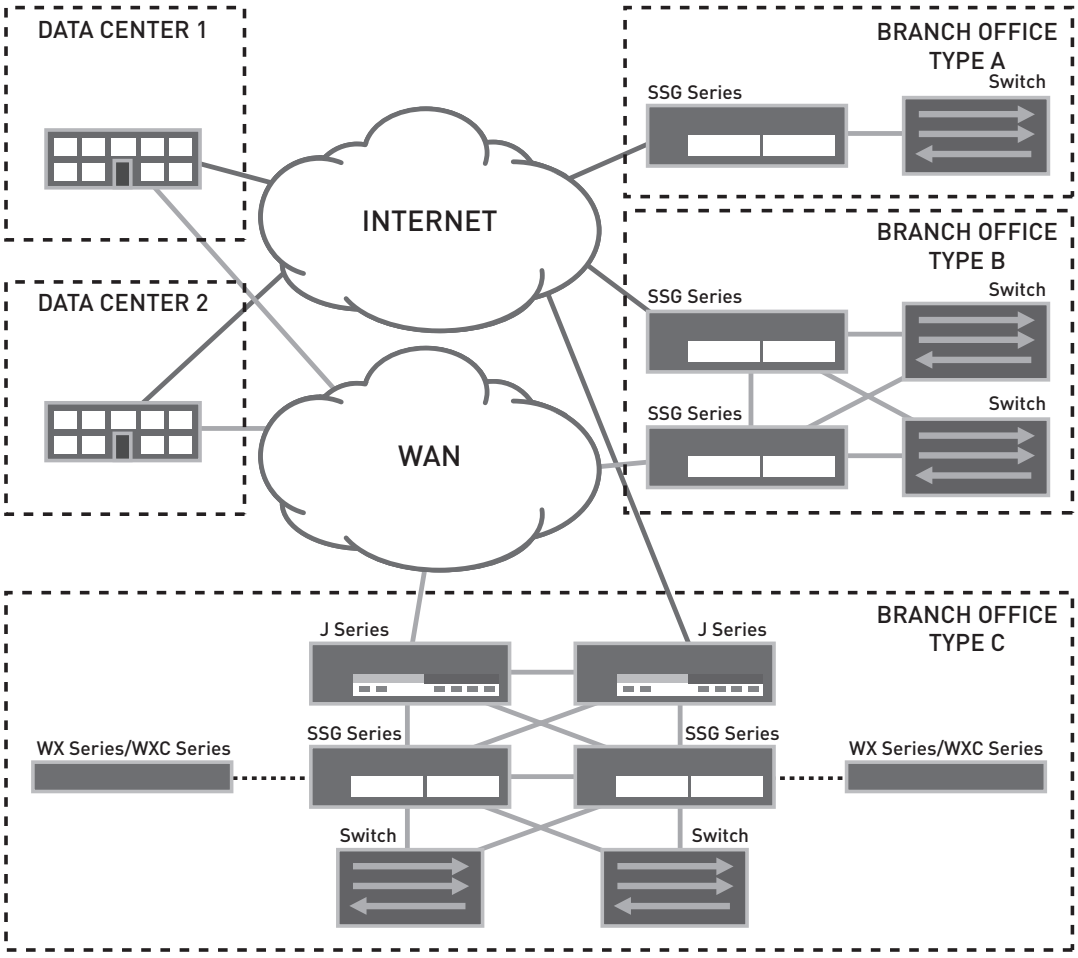

Figure 2: Branch office reference architecture

Copyright © 2010, Juniper Networks, Inc.

Table 1: Functionality, Features, and Capabilities of the Branch Office Types

Functionality	Feature	Capability	Type A	Type B	Type C
Security		Deep Inspection	•	•	•
		Antivirus	•	•	•
		Web Filtering	•	•	•
		Firewall	•	•	•
Connectivity	WAN	T1/E1	–	•	•
		MPLS	–	–	•
		Broadband	•	•	•
	LAN	Wired	•	•	•
		Wireless	Optional	Optional	•
High availability	Device Redundancy	–	•	•	•
	Link Redundancy	Optional	Optional	Optional	•
Performance optimization	WAN Acceleratiion	WAN Acceleration and Optimization	–	–	•
Manageability	Juniper Networks Junos® operating system	Single OS Site-wide Visibility and Control	•	•	•

Branch Office Connectivity over IPsec VPN

As a best practice, Juniper Networks recommends using an IPsec VPN implementation to provide branch office connectivity for the following major reasons:

- Satisfies standards-based compliance requirements with real-time and historical forensics and reporting
- Ensures network security at the branch office by deploying the same security solution that is at headquarters, optimized for the branch
- Offers data confidentiality and integrity via unparallel availability, agility, security, and manageability

The recommended solution for connecting large numbers of branch offices onto the enterprise network is employed through an IPsec-based VPN as a tunneling mechanism.

Design Recommendations

The following list summarizes the major design recommendations and their relationship to the branch types. Unless otherwise stated, these recommendations employ the following functionalities for the three branch office types:

- Scalability: The routing design should be scalable. A medium-sized number of sites (100 to 1,000 sites) must be deployable without significantly impacting the CPU resources or memory of the VPN devices.
- Redundancy: When redundant links are used, no single point of failure should exist. As depicted in the reference architecture, head and regional offices have more than one connection to the public and private networks. In the case of Type B and Type C branch offices, redundant paths must be provided to use all branch links to extend to the external networks.
- Route Summarization: Address summarization should be performed whenever possible, as a large number of remote sites are needed. Route summarization is important to reduce both the routing traffic and the memory consumed in the routers.
- Network Address Translation (NAT): Remote sites can be connected behind NAT devices. Some small branches and remote users share the Internet connectivity with other users, and in such cases, the IP addresses assigned to them can be private.
- Load Balancing: Load balancing of customer traffic should be achieved. Remote sites with dual-homing connections to the Internet are common. In those cases, optimum usage of both links is desirable while still providing a backup path whenever a link goes down.

Copyright © 2010, Juniper Networks, Inc.

- Configuration Simplicity: Provisioning is easy. When the number of sites is large, it is important to reduce the complexity of the configuration on a per-site basis, where possible. That is, requiring multiple configurations should be as simple as possible.
- Link-Failure Detection: Link-failure detection mechanisms are required. Because IPsec tunnels use a Layer 3 infrastructure, routing (and sometimes link routing) failures are possible where the tunnel termination points are not notified of the malfunction. Therefore, the control and data plane keep-alive protocols should be used to detect the malfunctions.
- Unified Threat Management (UTM): UTM features and firewalls should be enabled at the branches. For both to function properly, it is mandatory to guarantee that both ingress and egress traffic flows across the same firewall. The same is true for data centers. Traffic symmetry enforcement is required to ensure that each particular traffic flow (and sometimes groups of flows, as is the case with protocols that use more than one connection) is presented across the same set of inspection devices.

Routing Architecture

To ensure connectivity, a minimum of one IPsec tunnel between each branch office and the central office is required. With N remote locations, a full mesh of tunnels would equal N*(N-1)/2 total tunnels. In particular, if implementing a VPWAN solution (having 1,000 sites), this solution requires a minimum of 499,550 tunnels. Clearly, it would not only be quite difficult to manage but also requires each branch to terminate at least 999 tunnels.

Note: Enterprise networks in general have an asymmetric traffic profile—that is, most of the internal traffic travels from the branches to the data center and vice versa. Generally, traffic between branches does not constitute the bulk of the traffic carried across the network.

Tiered Network Approach—Layering the Network for Scalability

To circumvent the N2 scalability problem, it is customary to partition the network into several layers. Devices on each layer connect only to devices located on an upper and lower layer, as shown in Figure 3.

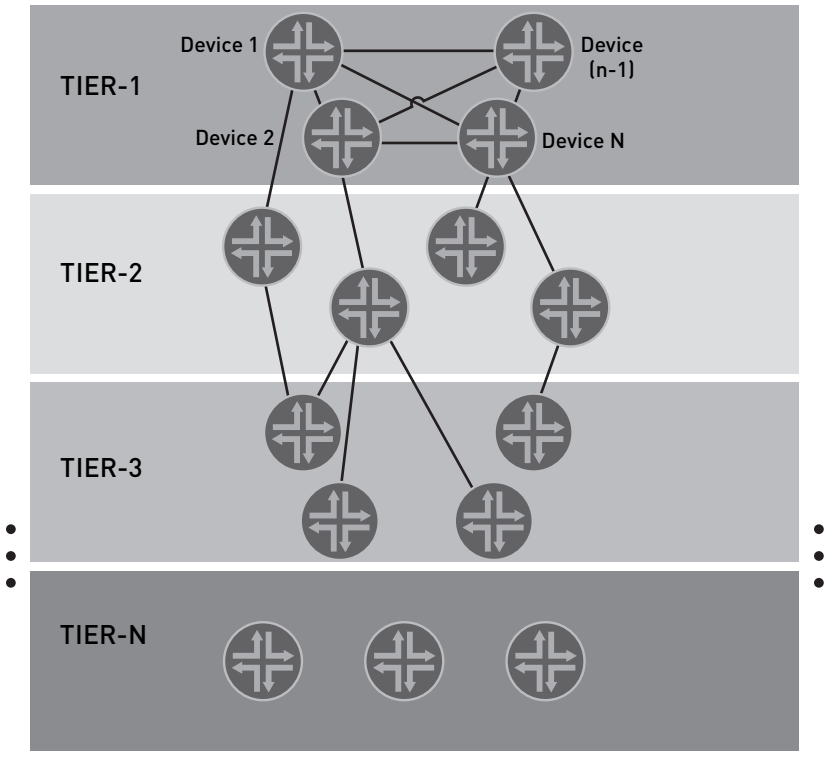

Figure 3: Multi-tiered/layered network architecture

Copyright © 2010, Juniper Networks, Inc.

When a device sends traffic to other devices, it sends the traffic to an upper layer router. The process repeats until the traffic arrives at the specified router with visibility to forward the traffic down to a lower layer. Because devices on the top layer are usually fully meshed, they can send the traffic to the appropriate router. The traffic is then forwarded down the layer chain until it reaches its final destination.

The obvious advantage of using a tiered layer approach is scalability. Any-to-any connectivity is required only for routers on the higher layer. Routers on lower layers require only a single connection (although for redundancy purposes it is common to provide more than one). This scalable approach reduces the minimum topology to a tree topology whereby each device requires only one or two connections (except for the devices on the top-most layer), effectively reducing the traditional scalability challenge to a linear function bound to the number of devices in the network.

Considering enterprises' scalability requirements, a two-tier design is sufficient to accommodate network needs. Because each VPN concentrator can terminate in excess of 2,000 tunnel connections, there is no need to further segment the design.

The resulting topology consists of a set of remote branches interconnected through a tier of data centers (or regional offices), as shown in Figure 4.

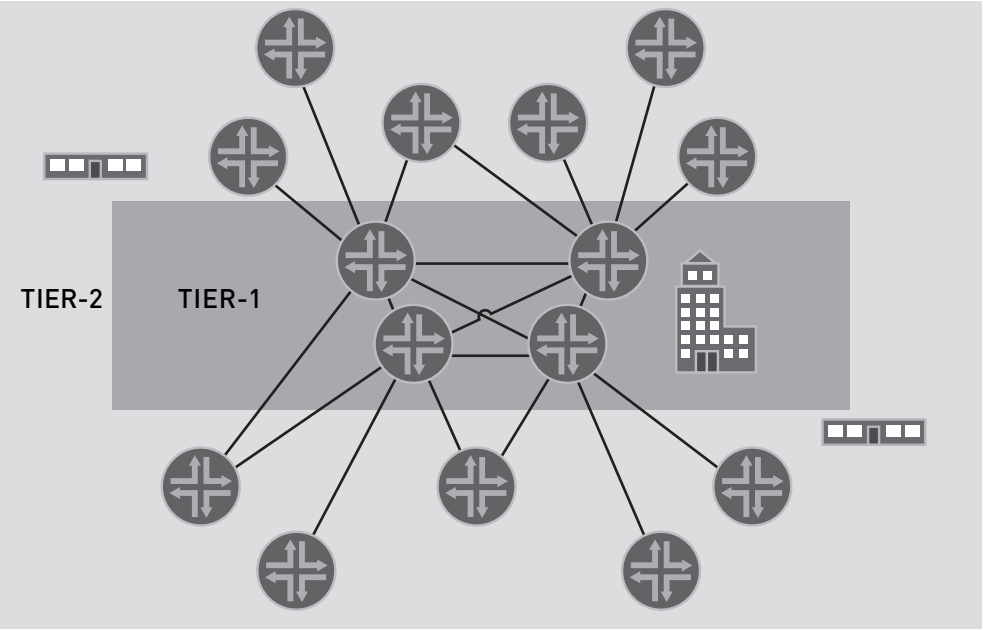

Figure 4: Two-tier network design for data centers

Note: Whenever traffic patterns require a hybrid design, Juniper Networks recommends using Auto Connect VPN (AC VPN). As an example, AC VPN is well suited if the group of branches has high demands of inter-branch traffic.

The AC VPN feature can ease the provisioning burden while providing full-mesh connectivity. For further details concerning AC VPN, see *Using AC VPN to Create Branch-to-Branch IPsec Tunnels.*

The proposed two-tier network design refers to the topology of the overlay IPsec network and not to the underlying physical network. Physically, the network resembles the diagram, as shown in Figure 2, Branch Office Reference Architecture.

Routing Information Protocol

Several designs for the routing topology were considered. On one side, a robust, scalable, and standards-based solution is desired. On the other side, ease of configuration and deployment is more appealing (considering the large number of sites expected). To address these concerns, different routing protocols (RIP, BGP, and OSPF) are discussed with their associated benefits and problems. For further details concerning BGP, see *Using BGP for Large Networks.*

Juniper Networks recommends RIP as the routing protocol for enterprise networks that connect between 100 to 1,000 branch office locations. Using a RIP-based IPsec routing implementation provides a cost-effective and secure alternative when employing leased lines and other expensive L2 networking technologies. However, some of the challenges as well as the trade-offs of each design decision are presented here. The proposed routing architecture's goal is to balance complex implementation with design scalability, but still be well-suited for small to medium-sized organizations.

Copyright © 2010, Juniper Networks, Inc.

Using RIP with On-Demand Circuits Extensions (Advantages and Disadvantages)

RIP is easier to provision and administer. When using RIP, each branch advertises the range of directly connected networks to each data center, using different metrics. Data centers advertise an aggregate of the network together with a more specific route to the networks that directly attach to the data center.

To reduce the amount of processing and traffic needs, demand circuit extensions should be enabled. With these extensions, route advertisements over a point-to-point connection (as in the case of the IPsec tunnels) are acknowledged and do not need to be periodically retransmitted, unless the network topology changes.

The main disadvantage of demand extensions is that they are not commonly used, and therefore are not supported by many routing vendors (including Junos OS-based routers). Integration with non-ScreenOS®-based devices might be difficult. However, a hybrid approach can be employed where non-conformant devices connect using standard RIP.

The advantages for using RIP with demand circuit extensions for a medium-sized network include:

- Low protocol overhead - Routing information is only exchanged in the event of failures or topology changes.
- Relatively fast convergence - VPN monitoring provides efficient and quick detection of end-to-end connectivity problems between two IPsec tunnel endpoints.
- Versatility - The routing path can be easily modified by using different metrics.
- Limited flooding of information - A link state database protocol, such as IS-IS or OSPF requires the flooding of routing information to all the devices in a single area (or level). This results in the constant redistribution of protocol information whenever a single device (or group of devices) fails. Because the network can have as many as 1,000 devices, many situations can occur where some of the IPsec tunnels could flap due to poor network conditions. In such cases, even when partitioning the network in several areas, the routing instabilities would be propagated throughout the network.
- Protocol simplicity and ease of operation - Many enterprise networks already use RIP as their preferred.
- IGP protocol.

Traffic Load Balancing for Type B and Type C Branch Deployments

Traffic should be balanced across more than one link when implementing Type B and Type C branch office topologies. Except for special cases, 1 10 Mbps link tends to be more expensive than 10 1 Mbps DSL lines. For reliability reasons, it is common practice to provide redundant links. Even when the driving force is to install an extra link because a backup connection is required—and because in North America customers generally pay a flat fee per link, independently of the traffic carried—it is desirable to use the extra capacity efficiently.

Solving traffic balancing at the data center is different when compared to the remote branches. Due to the large number of tunnels terminating at each data center, traffic can be balanced by splitting the tunnels and routing each group through a different link. If the traffic pattern is relatively dispersed among the different tunnels, this strategy provides excellent usage for every data center link.

However, at remote branches, it is *not* always possible to balance traffic in this manner. If the characteristics of the customer traffic were such that traffic could be easily distributed across multiple data centers, then simply routing the IPsec tunnels—going to the central offices through different links—would suffice.

When most of the traffic is directed to a single data center, traffic symmetry must be preserved. As with any traffic inspection device, the firewalls inspecting traffic (located at each head-end of the tunnels) must inspect traffic in both directions. Accordingly, for any particular connection, traffic going in and out of a particular data center must go through the same firewall.

Copyright © 2010, Juniper Networks, Inc.

Implementing ECMP as an Aid in Load Balancing

As shown in Figure 5, it is still possible to use both access links. Creating a pair of tunnels, each routed through a different connection (both originating on the same data center and terminating on the branch), provides the required load balancing. In this way, traffic is split using equal-cost multipath (ECMP) across both access interfaces.

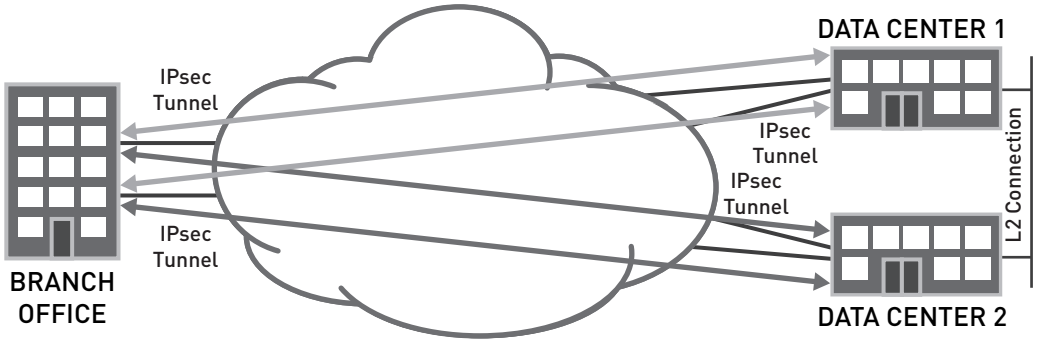

Figure 5: Branch with dual internet connections (load balancing using ECMP)

Finally, for the solution to provide both load balancing and data center redundancy, each branch office requires a set of *four* IPsec tunnels. As a rule, only two of these tunnels carry traffic to each data center, while the other two are dedicated to one of the data centers that experiences complete failure.

Load-Balancing Solutions in Relationship to Branch Connection Types

By observing the aforementioned considerations, there are three possible tunnel scenarios and their relationship to connection:

- Branch Offices with a Single Connection: These are single-homed branches that have a unique tunnel to each head office. Because each head office has two Internet links, the tunnels are evenly split between these two for a total of 500 tunnels per link.
- Branch Offices with a Single Connection to the Internet and a Single Connection to a Private Network (or PTP network): Branch offices connecting to the private network have a single tunnel to each data center using the private network. They also have a single tunnel to each data center through the public network. The tunnels across the private network are always preferred to those that use the public network. Therefore, there is no traffic load balancing in this scenario.
- Branch Offices with Dual Internet Connections: For branch offices with dual Internet connections that have two tunnels to each data center, ECMP provides load balancing for traffic across the tunnels.

Using Border Gateway Protocol for Large Networks

For large network deployments (more than 1,000 branches), Border Gateway Protocol (BGP) should be implemented because it is better suited to control route instabilities and a large number of network advertisements.

Using Border Gateway Protocol with Route Reflectors

The last design considered attempts to off-load most of the routing computations to a device other than the VPN devices. To do this, BGP helps propagate routing information with the aid of a route reflector, as shown in Figure 6.

Branch devices establish their IPsec tunnels with the central (or regional) offices. Each IPsec tunnel carries a single BGP session from the branch office to a route reflector located in the central office terminating the tunnel. In turn, the route reflector performs the route selection and sends the routing information to the VPN concentrator by using a single BGP session.

Copyright © 2010, Juniper Networks, Inc.

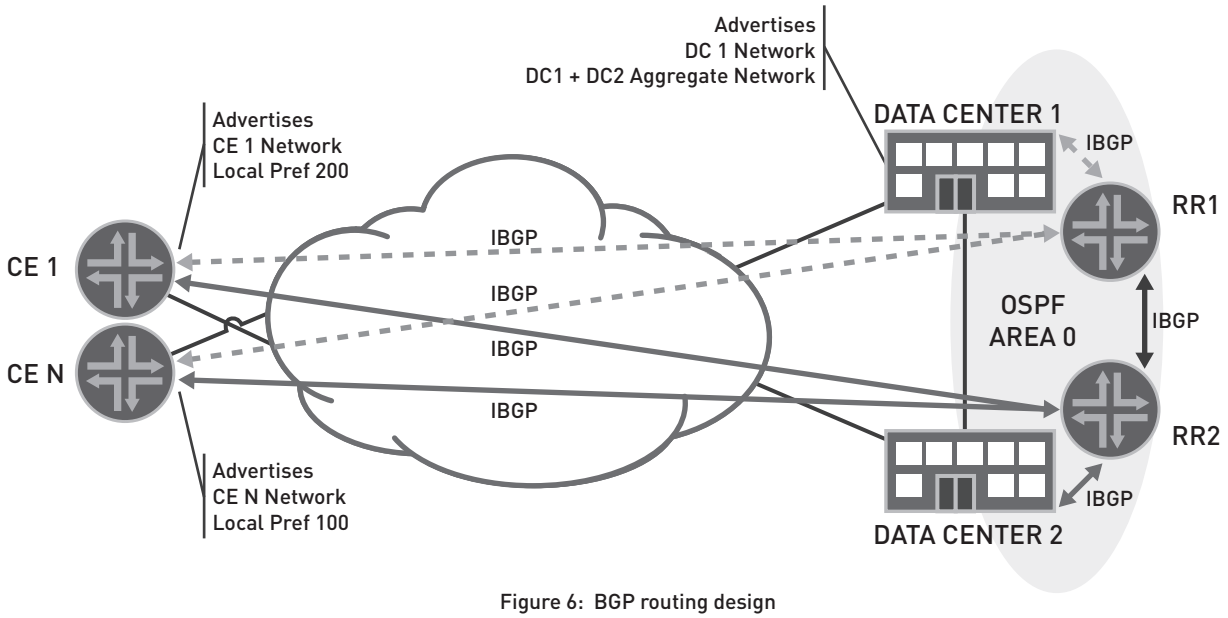

Figure 6: BGP routing design

The main advantages of this design are that it accommodates multiple devices and scales to a large number of remote offices. Route processing is somewhat distributed by route reflectors. However, each device still has to go through the BGP route selection process, but the number of sessions that each firewall has to maintain is minimal—as is the number of messages that it has to process.

Unfortunately, while ubiquitous on service provider environments, BGP is *not* commonly used by enterprise customers. This lack of expertise can prove to be challenging as the administration of the network becomes more complex. Also, the extra cost of the equipment—when employing route reflectors—must be weighed when selecting a final design.

Typically, when service providers sell the large-scale IPsec VPN service to their customers, they use BGP as the routing protocol in a similar manner. This solution is well tested and has been shown to work in large customer deployments.

Also, BGP might be a good choice as an interior gateway protocol (IGP) inside each data center (and between data centers). For information about using this protocol, see the *Internet Connectivity* section.

Using Static Routes

A simple solution used for connectivity consists of using static routes at both endpoints of the tunnel. On each branch, a single aggregate route for the entire internal network (and a more specific route pointing to the data center terminating the tunnel) is configured. In turn, at the data center, a route for each remote network is configured by mapping traffic to the particular tunnel. By modifying the metrics for the routes on both endpoints of each backup tunnel, traffic is directed to the backup tunnels only during a failure. An IGP routing—for example, OSPF—can then be used to distribute the static routes configured at each VPN concentrator. For additional information about the data center IGPs, see the *Internet Connectivity* and the *Internet Connections* sections.

Although basic to deploy, using static routes has several disadvantages:

- Provisioning and managing the routes can become troublesome particularly because each site can have from one to four tunnels. A minimum of 3,000 static routes must be configured (one on the data center for each site, and two at each site).
- Modifying the addressing space at a branch requires manual reconfiguration. The firewalls at the head-end of the network, terminating the IPsec tunnels, require reconfiguration with static routes pointing to the new addresses.
- Relying completely on an external form of dead peer detection (DPD)—such as IPsec DPD, Internet Key Exchange (IKE) keepalives, or VPN monitoring—is desired so traffic can be switched during a failure.

With this design, traffic originating at a device directly connected to a VPN concentrator and terminating a backup tunnel cannot reach the remote office associated with that tunnel. The problem resides on the protocol

Copyright © 2010, Juniper Networks, Inc.

administrative distances. Whenever the primary tunnel is active, traffic should be forwarded to the data center terminating that tunnel. Because static routes have low administrative distance, traffic forwarded by the backup device always chooses the local tunnel regardless of the metrics used. The return traffic, though, still uses the active tunnel, causing problems on the firewalls. This problem can be avoided by either using a router to commute the traffic to the appropriate VPN terminator or by changing the administrative distance of the static routes.

A variation of this design uses IPsec with aggressive mode and framed routes. In this scenario, each branch connects only to the primary concentrator. When connecting, xauth is used to authenticate the remote site and a radius profile is retrieved. The profile (using the framed-route attribute) indicates that the networks are configured on the branch, and the routes pointing to them are automatically provisioned. Whenever a problem is detected, both the central office and the branch clear the tunnels and remove the routes. The branch finally establishes a new tunnel using the backup terminator, and new routes are installed upon successful negotiation of the tunnel. Also, note that Junos OS and ScreenOS do not currently support framed routes.

This solution, while reducing the number of simultaneous active tunnels in the network, simplifies provisioning. Although it still relies heavily on DPD techniques, it avoids most of the drawbacks associated with static routes. Whenever the failover times are *not* critical, using single active tunnels and framed routes can provide an easily managed solution.

Due to the scale of the networks targeted, using static routes is quite difficult to manage, considering that load-balancing traffic of up to four tunnels per branch can be configured. This means that to use static routes, a special provisioning tool would have to be developed. Although convenient in some particular cases, it is impractical when designing a network for a non-homogeneous set of customers.

Using OSPF for Small Number of Branch Offices

Juniper Networks recommends using OSPF for networks that have less than *100* branch offices (locations), and only if the data center already uses OSPF for branch office-to-data center communication. The overhead associated with OSPF creates excessive C0PU101 utilization at the VPN device, especially when numerous branch offices are being connected. Using a link-state database protocol can cause routing information floods to all the devices in the same area. This results in the constant redistribution of protocol information whenever a single device (or group of devices) fails. Because the network could have a maximum of **1,000** devices, many situations can be found where some of the IPsec tunnels could flap due to poor network conditions. Even when partitioning the network in several areas, the routing instabilities would be propagated to all the routers in the same area.

However, it is possible to partition the network into multiple zones. In this case, instabilities are contained inside their respective zones, but routing information between areas is distributed in a way similar to how a distance-vector protocol operates. For these simple hub-and-spoke topologies, no major benefits result when using distance-vector-based routing protocols. Therefore, Juniper Networks recommends using either RIP or BGP as the preferred routing protocol for branch office-to-data center communications.

Additional Design Considerations

This section discusses some of the remaining design considerations.

- There is a maximum of **256** RIP neighbors allowed per interface. When implementing networks with more than 256 branches, several tunnel interfaces must be used.
- Demand extensions require an external neighbor monitoring mechanism, as explained in the ***RIP*** section. Current versions of ScreenOS can only use vpnmonitor for this DPD, which will be modified in future releases to notify the routing protocols when a particular IPsec neighbor is unavailable. As in most monitoring protocols, there is a trade-off between scalability and convergence speed. The design has been tested using a pooling interval of 2 seconds with a threshold of 5 seconds (providing a detection time in the order of 10 seconds).
- It should be noted that when RIP is administratively disabled with demand extensions, remote RIP neighbors will not be notified of the change, leaving stale prefixes in the routing table. However, this disadvantage does have a workaround by flapping RIP on the remote end.
- There is always the risk that during routing protocol convergence—when ECP is used over IPsec tunnels—forwarded traffic uses one IPsec tunnel while the return traffic is received on a different tunnel. By default, a ScreenOS-based firewall would drop this traffic (as it would fail the RPF check). In addition, the command set flow reverse-route tunnelprefer allows packets to be received on a different IPsec tunnel during such times.

Copyright © 2010, Juniper Networks, Inc.

- Assumptions about the nature of the IPsec tunnels have not been made. Both aggressive and main mode tunnels can be mixed on the presented network. If a failure occurs, when using the aggressive mode with dynamic peers, only the remote peer can initiate a new IPsec tunnel connection. This might result in longer recovery times.
- It is important to configure RIP using demand circuit extensions. Otherwise, the protocol overhead would be large, as every branch office would retransmit its prefixes during every update interval (30 seconds by default).

For detailed instructions about using a RIP-based IPsec routing implementation, refer to the *Implementing IPsec VPN for Branch Office Connectivity Using RIP. Appendix C* lists the product tables for the various Juniper Networks and Juniper partner product solutions that support the design and deployment of high-performance enterprise networks.

Using Auto Connect VPN to Create Branch-to-Branch IPsec Tunnels

Juniper Networks developed the Auto Connect VPN (AC VPN) feature, which allows the dynamic creation of branch-to-branch IPsec tunnels. These tunnels are created on an on-demand basis and are triggered by the traffic generated at any given branch office. To accomplish this, AC VPN uses Next Hop Resolution Protocol (NHRP). This protocol was originally developed for non-broadcast multiple access (NBMA) networks and is intended to provide a discovery mechanism for stations to determine the L2 address of a device that connects to a particular L3 network (or the egress router for that particular destination).

Next Hop Resolution Protocol

NHRP is reused and augmented to achieve a similar task—that is, to discover the public IP address of a VPN termination endpoint. Whenever a branch office needs to send traffic to another branch office, this office can establish an IPsec tunnel directly to the destination branch. To this effect, the branch originating the traffic can use NHRP to discover the public IP address of the remote branch.

Some proprietary extensions have been added to the protocol and provide a way to simplify the provisioning of these tunnels. Before presenting the details, it is important to understand the required base topology of the network for compatibility with NHRP.

For AC VPN to work, it is necessary to have a star topology network that connects all the branches to a central hub, as shown in Figure 7. The branch offices use these tunnels to register the networks directly connected to each of them. The regional office stores (in a local database) a mapping of all the networks that each branch office registered, together with the public IP address that each branch uses to terminate IPsec. Some additional information that helps the branches to authenticate each other also is stored in the local database.

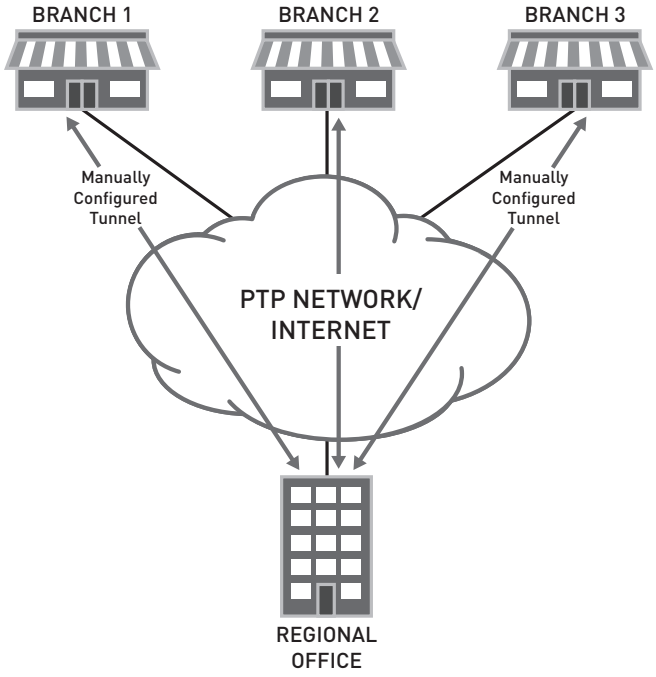

Figure 7: Star topology – connecting branches to central hub

Copyright © 2010, Juniper Networks, Inc.

Note: The hub also stores a profile with the configuration of the IPsec tunnels that branch offices use to achieve connectivity. This way, the configuration is simplified, as the tunnels only have to be configured on the hub. This configuration is then pushed to the spokes whenever a direct IPsec VPN connection is established.

The Tunnel-Building Process

Once the registration process is finalized, the branch offices start building tunnels (Figure 8) between themselves as follows:

- A branch office has traffic to send to another branch office. Normal IP routing occurs and the traffic is sent to the hub so that this traffic can then be forwarded to the destination branch.
- The hub VPN concentrator forwards the packet and it notifies the NHRP module that there is traffic going across the hub from two networks that have mappings stored in the next-hop server (NHS) cache.
- The hub concentrator then sends an NHRP resolution packet to the branch, along with a mapping of the remote branch office network, to its public IP address. In addition, the hub concentrator sends a hash of the certificate that the remote branch uses to identify itself, and a profile describing the configuration of the IPsec tunnel that each branch office should use.
- **Note:** This information is encrypted over the IPsec tunnels (established between the hub and spokes) so the trust relationship has already been determined.
- The branch can update its NHRP cache information after receiving the mapping, and using this information, establishes a tunnel to the remote branch.
- The two branches—after the tunnels have been established—add a route to the other's branch network through the newly created tunnel. These new tunnels are tagged as NHRP routes.

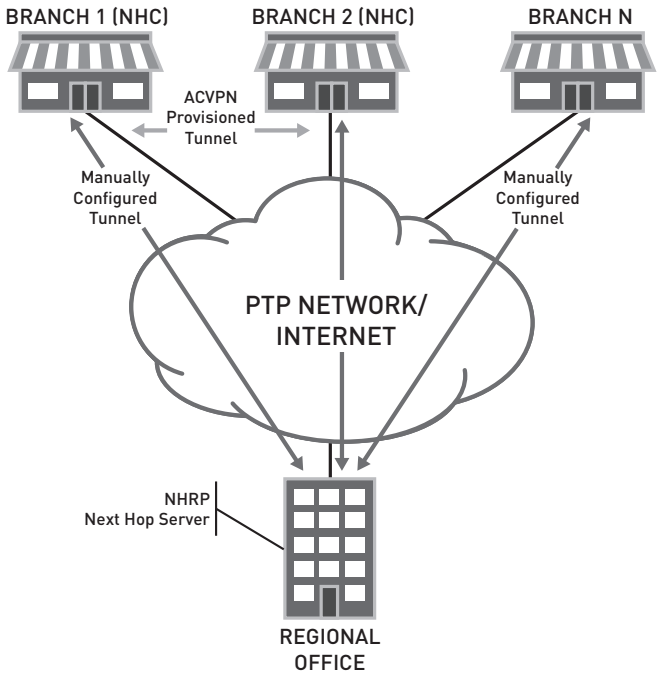

Figure 8: AC VPN provisioned tunnels between branches in the same region

Auto Connect VPN Design Considerations

The following are the design considerations and assumptions associated with this implementation.

- The NHS address must be the address of the tunnel interface terminating the IPsec tunnels from the branch offices. In particular, the NHS will not detect requests on loopback interfaces.
- A device can only act as a next-hop client (NHC) or an NHS but *not* both because hierarchies are *not* supported.
- The Type B branch offices do not have AC VPN provided for the secondary device. During a failure, the AC VPN service is *not* available and traffic is routed through the hub.

Copyright © 2010, Juniper Networks, Inc.

- The security associations (SAs) and the NHRP caches are not synchronized when active/active NSRP is used. If a failover occurs, a new NHRP registration is performed, and as a result, branch-to-branch tunnels must be reestablished. However, reestablishment of tunnels will not impact the branch-to-branch traffic, as branch traffic still will be routed through the hub.
- The branch offices that are only connected to the same hub (that is, a data center or regional office) can establish IPsec shortcuts between themselves. When branches are not connected to the same regional office/ data center, traffic flows using the pre-existing topology.
- AC VPN establishes shortcuts only between branch offices connected to the same hub for multi-tier topologies, as illustrated in Figure 9. In the example network, only branch offices in the same region can establish shortcuts. However, traffic between branch offices still can use normal routing and go through the different hubs until the traffic reaches the desired destination.

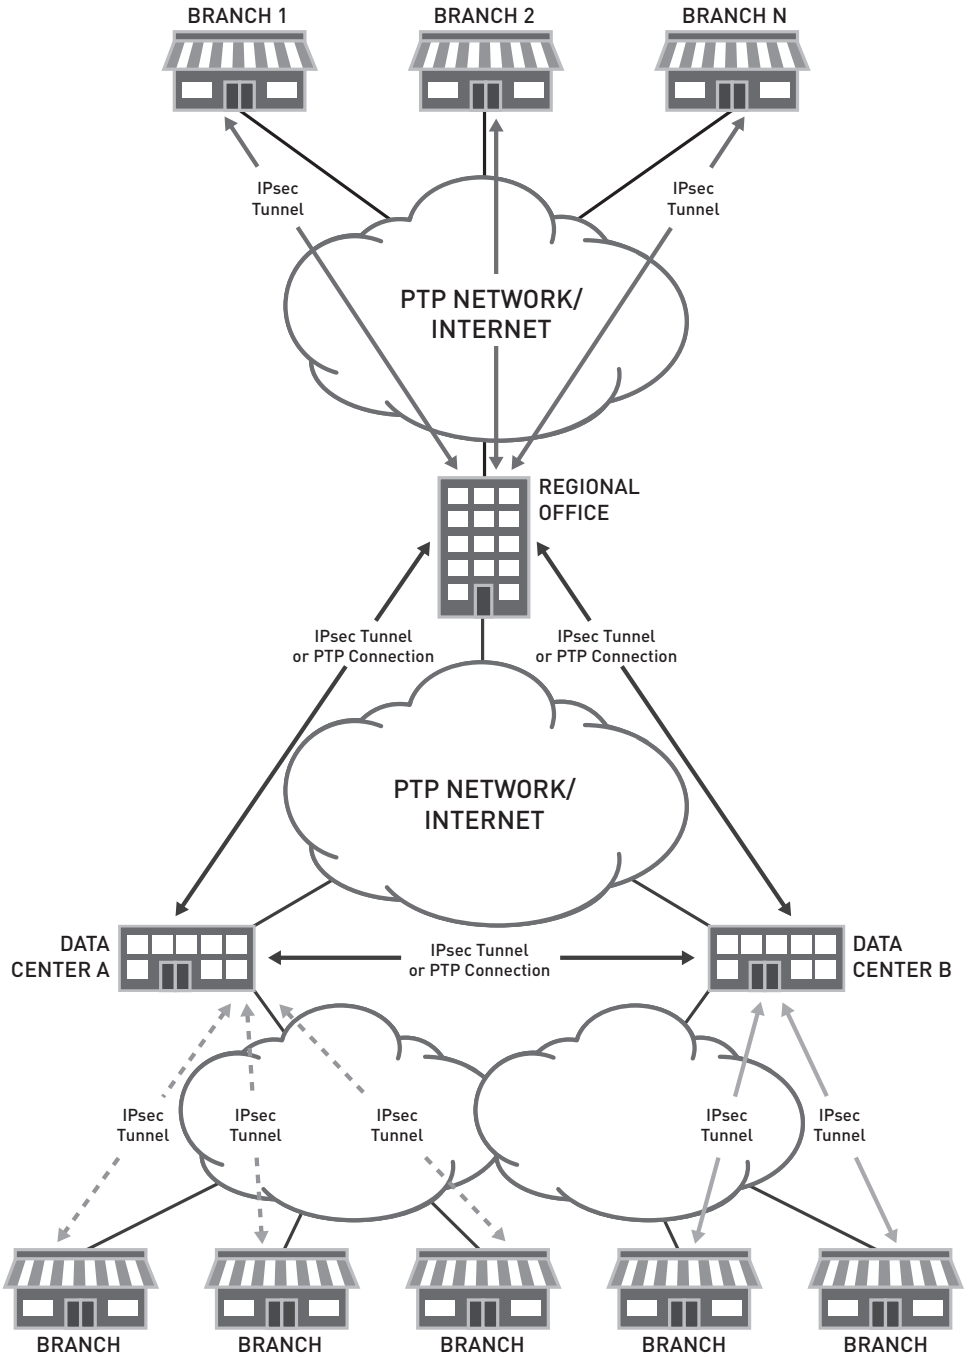

Figure 9: Multi-tier topology

Copyright © 2010, Juniper Networks, Inc.

- A single NHS server only can be configured on a per-client basis. During a complete failure at the hub (either data center or regional office acting as an NHS), branch offices cannot establish shortcuts until connectivity to the hub is restored.
- A new registration to the NHS is required when an NSRP failover is triggered. If a failover occurs at one of the hubs, then every branch office has to reregister and the NHRP cache has to be repopulated.
- The NHRP is not supported over unnumbered interfaces.

High Availability for the Branch Office

Branch office HA is a key design concern that assures business continuity for the branch offices. Juniper implements branch office HA by using link, device, or a combination of link and device redundancy to ensure network availability. Juniper offers three types of configurations, differing only by the branch office profiles. The result is a high availability enterprise network that can reliably connect the branch office locations to the data centers.

High Availability Requirement Levels (Link, Device, Device and Link Levels)

When defining HA, first you must identify the level of HA required for each branch office. The three levels of high availability include:

- Link-Level HA: This requires two links to operate in an active backup setting so if one link fails, the other takes over (or likely reinstates) the forwarding of traffic that has been previously forwarded over the failed link.
- Device-Level HA: This means effectively doubling up on devices to assure there is a backup device to take over for a failed device in such an event. Typically, the link redundancy and device redundancy are coupled, which effectively ties failures together.
- Device- and Link-Level HA: This allows a device to fail without requiring the respective link to fail. Note that there still will be a device attached to each link, and if that device fails, a link failure may occur as well. However, not every device failure will cause a link failure.

High Availability Functionalities

Juniper Networks Enterprise Framework and *Branch Office Reference Architecture* documents present the framework for this solution. Table 2 summarizes the key HA design functionalities that are used as a basis for branch office HA design recommendations.

Table 2: HA Functionalities

Functionality	Description
Link failure protection	Failure on any given access link should not result in connectivity loss. This only applies to branch offices with at least two upstream links connected either to a private network or to the Internet.
Device failure protection	No single device failure should result in connectivity loss from the branch office to the data centers (except for Type A branch offices, which do not provide redundant devices). However, a failure in a device might result in Internet connectivity loss if only one Internet connection is used.
Data center failure protection	In the event of a complete failure in one of the data centers, traffic must be rerouted to a backup data center, as data centers share a point-to-point connection.
Session persistence	Branch offices with redundant devices should provide session persistence. That is, in the event of a failure, established sessions should not be dropped, even if traffic was being forwarded by the failed device.
Load balancing	Customer traffic is balanced across dual connections to the data center. If a link fails, all traffic is directed through the remaining link.
Traffic symmetry	UTM features and firewalling are enabled at the branches. For these to work, this design guarantees that both ingress and egress traffic flows traverse the same firewall. The same scenario is implemented at the data centers.

Copyright © 2010, Juniper Networks, Inc.

Functionality	Description
Network Address Translation (NAT)	NAT is enabled so that machines in the trusted and guest zones can access the Internet. In the event of a failure, Internet sessions might not be preserved as the translated addresses of that traffic might have to change and different service providers might be used on the Internet links.
Fast failover times	Whenever failures occur (link, device, or data center failures), traffic should be rerouted in less than 30 seconds. Within this time, packet loss might occur, but sessions will be maintained if the user applications can withstand these failover times.
Secure management traffic	SNMP is enabled through the IPsec tunnels. Whenever backup devices are provided for branches Type B and Type C, it is possible to monitor both devices—even if one of them is not forwarding traffic (nor terminating any IPsec tunnel).

High Availability for Branch Office Type A

To implement HA, the best practice calls for each Type A branch office to employ two Internet connections and four tunnels (two per each data center). Traffic is load-balanced across each pair of tunnels. That is, whenever traffic is directed to a given data center, sessions are load-balanced in a round-robin fashion across each IPsec tunnel going to that data center. In turn, the tunnels are configured so that each tunnel uses a different egress link, resulting in a limited HA configuration for this branch office type.

Figure 11 shows as an example, branch office Type A HA configuration with dual Internet connections to a single data center. Note, that although multiple data centers might be used, from the branch HA perspective, the configuration is identical. Only the IPsec tunnel configurations and routing have changed. For further details, refer to the *Implementing IPsec VPN for Branch Office Connectivity Using RIP* application note.

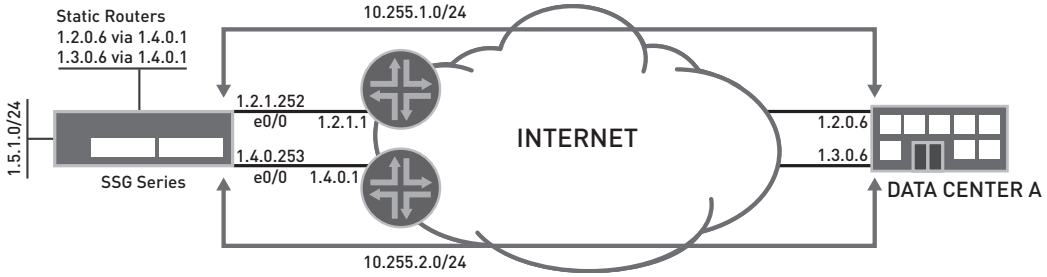

Figure 10: HA configuration for Type A

VPN Security Zone Configuration for Type A

Figure 11 shows the VPN security zone configuration. VPN tunnels are part of a separate zone named the "VPN zone." Implementing a VPN zone is an important consideration when designing security policies, as traffic going to the data centers (or other branches) will exit through this security zone.

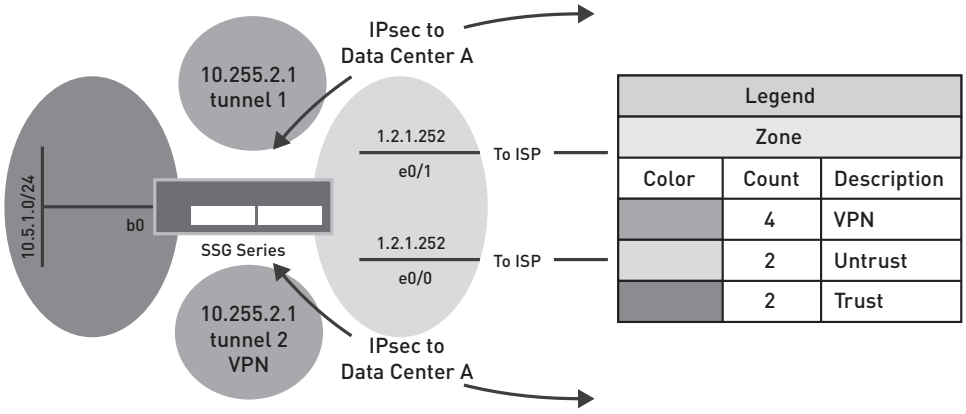

Figure 11: VPN security zone configuration for Type A

For Type A, NAT is provided based on the egress interface. That is, whenever traffic is routed through interface Ethernet 0/0, traffic is NATed using the interface's IP address as the source address (which is assigned by an Internet service provider (ISP) and possibly different than the one used in interface Ethernet 0/1). Similarly, whenever traffic

Copyright © 2010, Juniper Networks, Inc.

is routed through interface Ethernet 0/1, the interface IP address is used to NAT the traffic. In this way, there is no need to propagate the addresses between service providers. In addition, DSL connections are supported because it is not necessary to know the IP address assigned to each of the Internet-connected interfaces in advance. Instead, the Dynamic Host Configuration Protocol (DHCP) is used in this case.

For deploying HA at branch office Type A, see *Implementing High Availability (HA) at the Branch Office.*

High Availability for Branch Office Type B

The branch office Type B uses two firewall devices that are connected to two different networks, such as a Peer-to-Peer (PTP) network and the Internet. IPsec tunnels are configured to each data center using both networks in a similar fashion like branch office Type A. The difference is that the metrics are lower on the tunnel interfaces terminating the IPsec tunnels that transverse the PTP network. Thus, whenever possible, the PTP network carries traffic going to and from the data centers.

Figure 12 shows the HA configuration for branch office Type B.

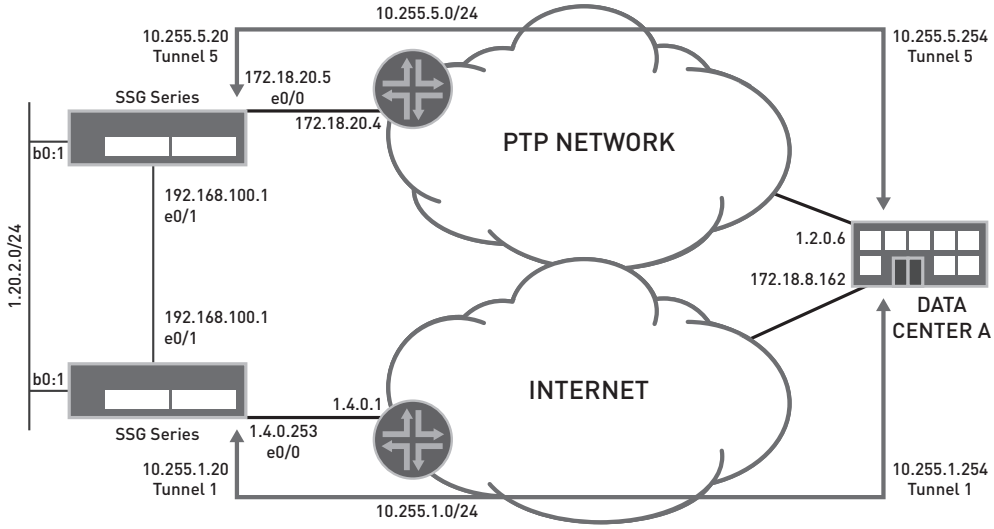

Figure 12: Type B optimized – HA configuration

Using Secure Services Gateway for Type B

For branch office Type B, each Juniper Networks SSG Series Secure Services Gateway terminates a pair of tunnels (one to each data center), as each is connected to a different network. Both devices are constantly active, but the NSRP is used in the trust (and guest) interfaces to direct the traffic to the SSG Series that connects to the PTP network. NSRP is configured in such a way that whenever a tunnel fails, NSRP fails over to the SSG Series terminating tunnels routed through the Internet. In this way, the PTP network is preferred over the Internet if the tunnels are active. Whenever a tunnel fails at any of the data centers, traffic is rerouted to the secondary SSG Series gateway.

Whenever the primary SSG Series is active, Internet traffic is routed using the Ethernet interface connecting both SSG Series Secure Services Gateways (belonging to the sync zone). Traffic, in turn, is NATed to the egress interface address on the ISP that connects the SSG Series (1.4.0.253 in this example).

The PTP network is used to back up the Internet connection whenever the link between the SSG Series and the Internet fails. One of the data centers advertises a default route over the PTP-transported IPsec tunnels. A default route is also advertised by the SSG Series to its neighbor over the Ethernet that is connecting them (Ethernet 0/1 in our example). In this manner, when the connection between the SSG Series and the Internet fails, the other SSG Series prefers the default route received through IPsec and sends all of its Internet traffic to the data center.

Because addresses in the PTP network are known in advance, a tunnel terminating at the primary SSG Series uses main mode and identifies the IPsec peers by their remote IP address. Instead, tunnels routed through the Internet use aggressive mode and IDs identify peers.

Copyright © 2010, Juniper Networks, Inc.

There is no session or configuration synchronization between the SSG Series Secure Services Gateways. Session persistence happens by disabling TCP SYN checks when flows are created inside IPsec tunnels. In this way, when traffic is rerouted to the secondary SSG Series, a new session is created and the traffic is forwarded to the destination. Note that this does *not* represent a security vulnerability point because SYN checks are enabled at the remote end of the tunnels.

It is important to understand that policies allowing traffic from the trust zone to the VPN zone (see Figure 13) have to be mirrored in the opposite direction (that is from VPN to the trust zone). When traffic is forwarded over a failover device, retry packets might originate only from the VPN zone to the trust zone (depending on the nature of the traffic). If policies do *not* allow this traffic, a new session is not created and traffic—from sessions created before the failure—is dropped.

Note that mixed mode NSRP is used in this configuration. Because Interface Bridge 0 has a Virtual Security Interface (VSI) (bridge 0:1) that is used as a default gateway, the egress interface used is either a tunnel interface (for VPN traffic) or a non-VSI Ethernet interface for Internet traffic.

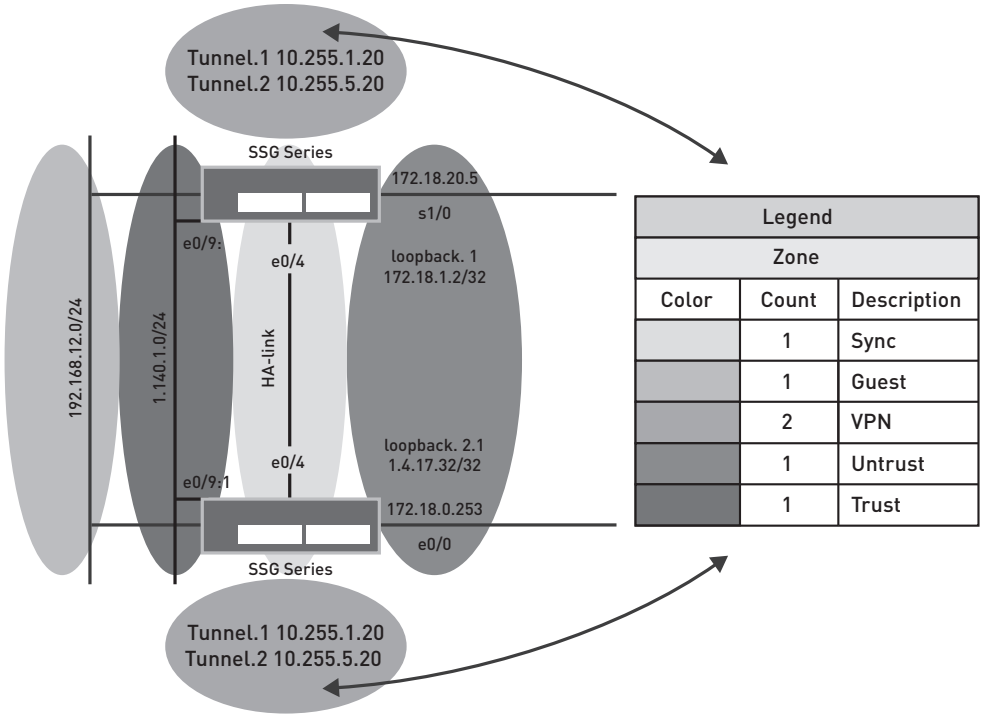

Figure 13: Type B – security zones

For HA deployment at branch office Type B, refer to *Implementing High Availability at the Branch Office.*

High Availabilty for Branch Office Type C

Medium- to large-sized branch offices use the Type C branch office configuration. This type of configuration provides session redundancy with session replication, as well as traffic load balancing when multiple connections to the same network are used—for example, multiple Internet connections. Figure 14 shows the HA configuration for branch office Type C.

Copyright © 2010, Juniper Networks, Inc.

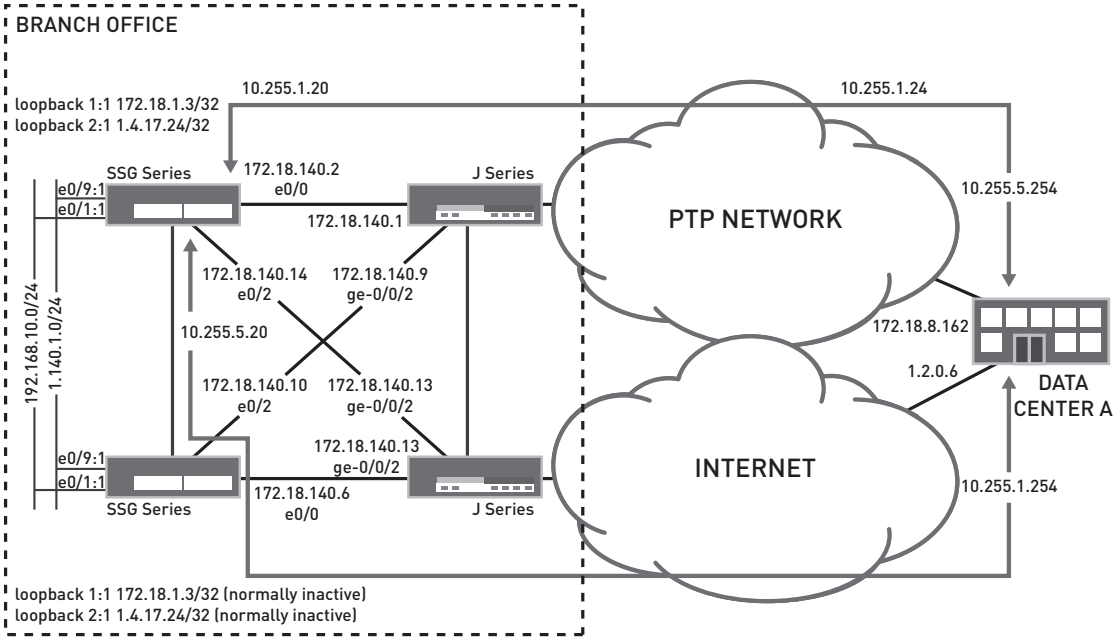

Figure 14: Type C – HA configuration

Branch office Type C uses OSPF to advertise loopback interfaces used to terminate the IPsec tunnels. The NSRP that is monitoring the interfaces facing the trust and guest zones (as well as the interfaces connecting to the Juniper Networks J Series Services Routers) determines the status of this loopback interface. The loopback interfaces terminating the tunnels are the VSIs that are part of Virtual Security Device (VSD) group 1. Whichever device has this VSI active—SSG Series (A) or SSG Series (B)—propagates the VSI addresses to the J Series routers using OSPF. Similarly, the J Series (B) router injects a default route into OSPF, while the J Series (A) router injects a route to the PTP network that also uses OSPF. Figure 15 illustrates this process.

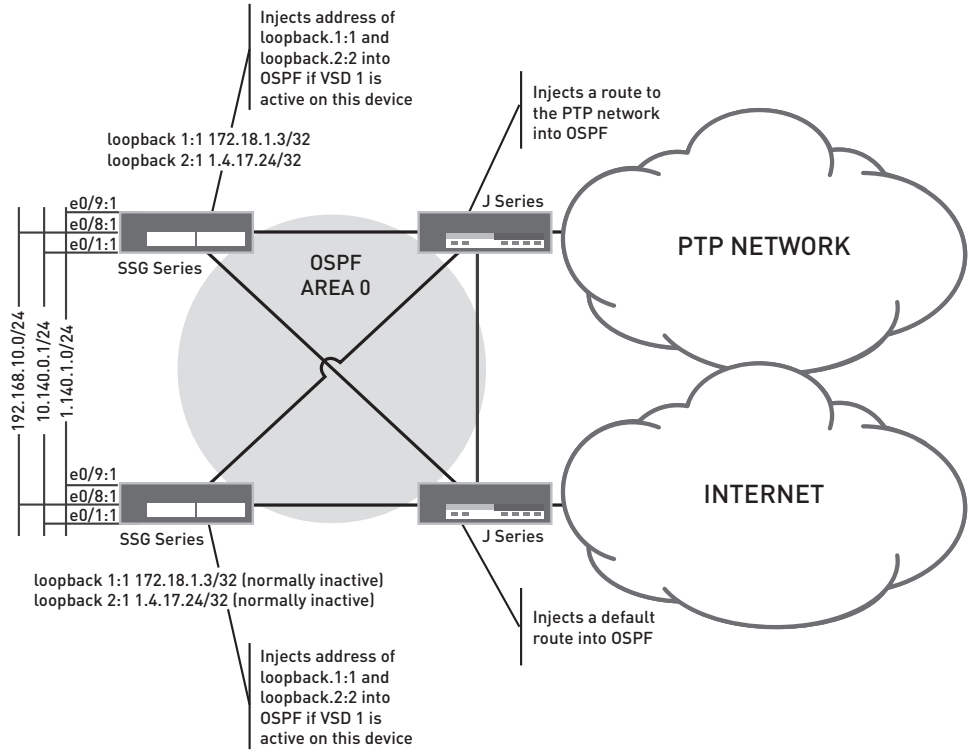

Figure 15: Intra-branch using OSPF

Copyright © 2010, Juniper Networks, Inc.

Pertaining to the other branch offices, the interfaces facing the J Series routers and the loopback interfaces are part of the untrust zone, the tunnel interfaces are part of the VPN zone, and the guest and user-facing interfaces are part of the guest and trust zones, respectively. See Figure 16.

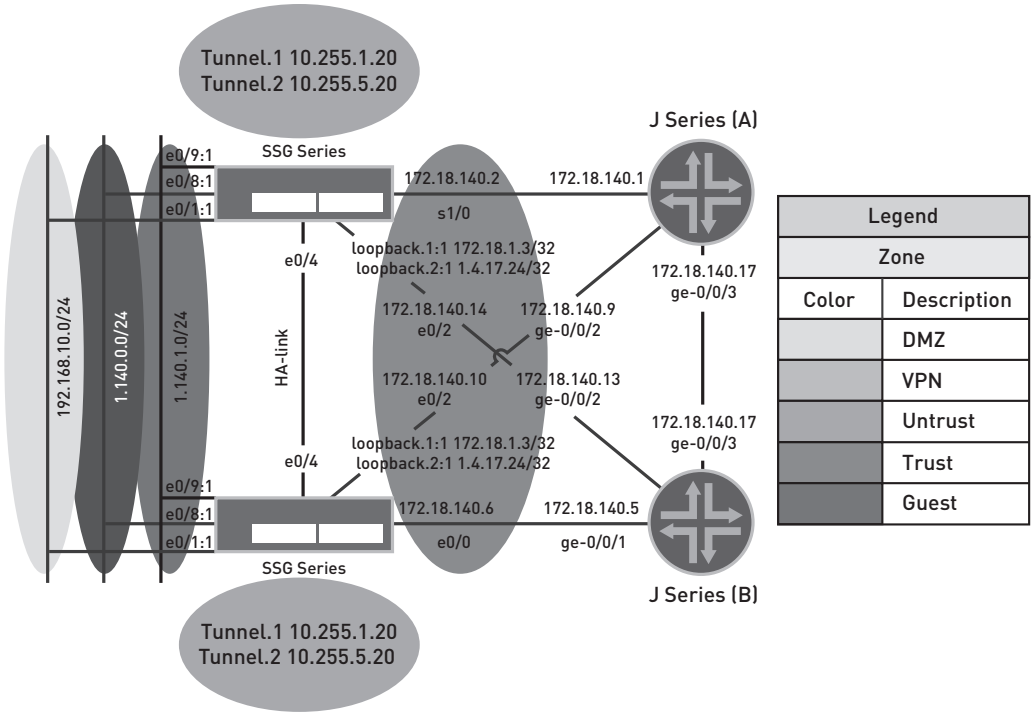

Figure 16: Branch Type C – security zones

Under typical network conditions, IPsec tunnels terminate on the SSG Series (A). In the event of a failure, the VSIs on the SSG Series (B) become active (for instance, if the interface Ethernet 0/9:1 that connects to the trust zone fails). The IPsec tunnels then terminate on the SSG Series (B) because OSPF advertises this device as hosting the VSI interfaces. The configuration of the SSG Series (B) is similar to the SSG Series (A), except for the NSRP priorities for the VSD group they share—where the SSG Series (A) is preferred—and their interface addresses.

For deploying HA at branch office Type C, see *Implementing High Availability at the Branch Office*. In addition, *Appendix C* presents the product tables for the various Juniper Networks and Juniper partner product solutions that support the design and deployment of high-performance enterprise networks.

Connectivity at the Data Center

Accommodating secure branch office connectivity requires a corresponding data center solution that offers HA, supports IPsec VPN connectivity, and delivers assured network performance and manageability across the enterprise. The data center design must support branch office connectivity while ensuring business continuity for the entire enterprise.

Implementing a High Availability Enterprise Network at the Data Center

This section offers design guidance and practices for implementing a resilient central edge design that contains no single point of failure and consists of a fully meshed, redundant active/active high availability deployment. This section also provides Juniper's design practices and recommendations for creating an agile and highly available network at the data center. While this design primarily focuses on the firewall deployment, it also describes other elements needed to design a highly available network infrastructure at the data center.

Design Requirements

Juniper Networks *Enterprise Framework and Branch Office Reference Architecture* documents provide the framework for this solution. Table 3 summarizes the key design requirements. The requirements are grouped by functionality, which further describes the design criteria for implementing this high availability network deployment.

Copyright © 2010, Juniper Networks, Inc.

Table 3: Data Center Key Design Considerations

Requirements	Description
Internet connectivity	• The design must employ a minimum of two Internet links. • The edge-connecting routers must provide redundancy as well as ensure service accessibility. • The active/active Internet connection requires two edge routers to provide resilient Internet connectivity. • A BGP feed is required from each of the providers to enable failover. • A rate limiting of traffic to the firewall is needed so that a flood of traffic from the Internet does not affect the network. • A stateless inspection or packet filtering must be used.
Private WAN	• Private circuits must be either point-to-point connections or connect over a provider-provisioned MPLS network. • All traffic that originates from the branch that is destined for the data center must be encrypted. • Private WAN is deployed off of the VPN firewalls.
Firewalls	• Internet firewalls must host the network operations center (NOC). • Firewalls must connect to the Internet and receive routing information from the Internet edge routers. • IPsec VPN firewalls provide the connectivity hub for all remote sites and they terminate IPsec VPNs from the Internet as well as from private WANs. • The IPsec firewalls must terminate VPN tunnels for all of the remote branches over the private WAN. • The following must be employed: redundant hardware, dynamic routing protocols (DRP), and fully meshed links. • The design must allow for a highly scalable VPN services infrastructure without being dependent on the availability of Internet firewalls.
Shared services	• The Internet firewalls must have a default route (obtained from the Internet edge routers) into the shared services core. • The connectivity to the firewalls must be in a mesh deployment. • The routing on the shared services core must integrate with the firewalls.
High availability	• The design must use a meshed solution to provide redundant paths on each redundant device. Internet connectivity. • The design must employ a minimum of two Internet links. • The edge-connecting routers must provide redundancy as well as ensure service accessibility. • The active/active Internet connection requires two edge routers to provide resilient Internet connectivity. • A BGP feed is required from each of the providers to enable failover. • A rate limiting of traffic to the firewall is needed so that a flood of traffic from the Internet does not affect the network. • A stateless inspection or packet filtering must be used.

Copyright © 2010, Juniper Networks, Inc.

Data Center Network Architecture

Figure 17 illustrates the data center network architecture. The design employs a redundant network topology so that user traffic continues to be forwarded despite failures of one or more links or nodes in the network. This implementation includes the following key features:

- Provides dynamic routing protocols
- Offers a fully meshed interface deployment
- Contains stateful failover

The architecture is derived from a working and tested example of the configuration and provides the basis for the design practices and information presented in this guide.

At the WAN edge, network architects have often used L2 switches to form a hierarchical mesh so that the multitude of links provides fault protection in case of failure. The design presented in this section employs Juniper Networks M Series Multiservice Edge Routers and SSG Series firewalls, and leverages the routing functionality of the SSG Series to provide a routed connectivity solution instead of a traditional switched mesh. Using this design places failure detection and correction into a domain that is solely routed, providing more effective and intelligent uses of network resources. The direct protocol interaction between the routers (without intervening switches) eliminates the typical layer of Ethernet switches commonly used at the edge.

The design uses OSPF as the interior gateway protocol between the security gateways and edge routers and uses a mixed mode NSRP to ensure that hosts can always reach the routers. The firewalls are seamlessly integrated into the routing domain because if there is a topology change, OSPF dynamically changes the forwarding path from the primary to the secondary firewall. OSPF link costs control routing paths in a deterministic manner, which eliminates the possibility of asymmetric routing. OSPF manages path calculation through the network topology and advertises routes between the WAN routers and the internal network.

Because the design uses fewer nodes, it reduces troubleshooting errors and the number of potential failure points. This design moves the security devices closer to the provider edge and decreases the number of devices that can be compromised due to hacking.

The details of the design are discussed as follows:

- Internet Connectivity
- Firewalls (Internet and VPN)
- Shared Services (Core)
- High Availability

Copyright © 2010, Juniper Networks, Inc.

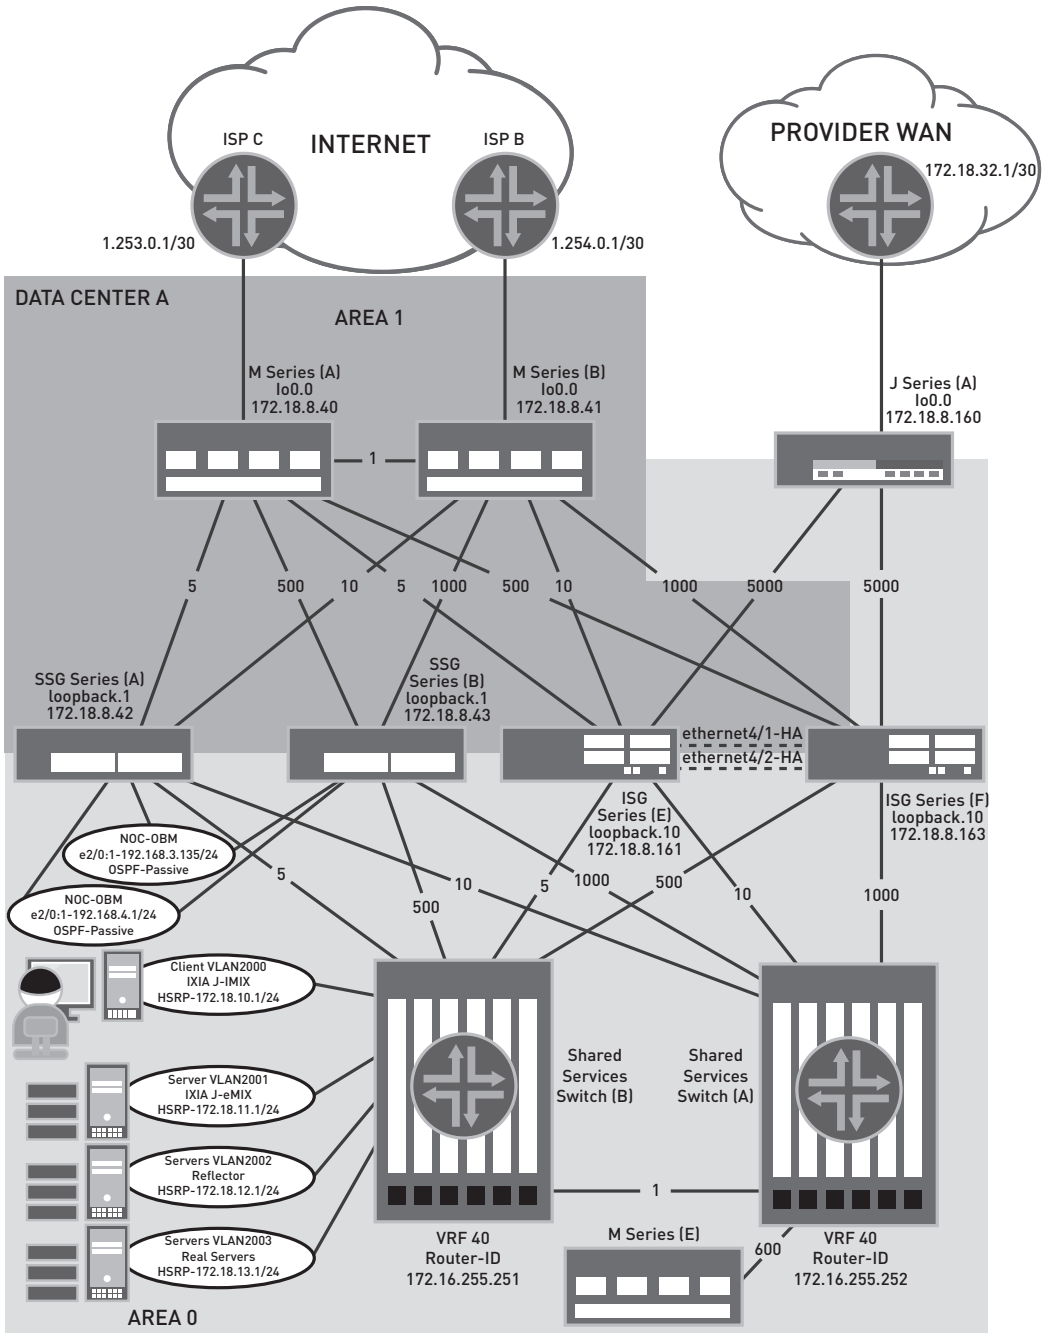

Figure 17: Enterprise network for the data center

Internet Connectivity

The Internet connectivity design (Figure 18 and Figure 19) consists of the following major components:

- Internet Connections
- BGP/EBGP
- Edge Routers

Copyright © 2010, Juniper Networks, Inc.

Internet Connections

Because HA is an integral part of the design, the solution uses two provider links. The active/active Internet connection uses two edge routers, which provides device redundancy and ensures service accessibility. BGP feeds provide failover protection from each of the providers, allowing the Internet to use the best path to the local network in the event of failure. Routing information is passed back to the network core using the IGP.

Because the design uses two separate Internet links, you can connect to two completely separate networks. Each link is diverse and provides a different peering design on the Internet. It is possible to use more than two different providers—that design decision is left to the discretion of each organization. The more connections provided potentially increase the availability of the Internet.

Border Gateway Protocol or External Border Gateway Protocol Operation

Each router has one external BGP or EBGP session for connecting directly to the provider. To give the design routing continuity to the Internet, each of these routers has what is called an "internal BGP session" to each other. Internal BGP is a session that occurs within one autonomous system, whereas EBGP occurs between two autonomous systems. This allows the network to determine the best possible path to the remote network.

The goal here is *not* to perform load sharing, but to provide the most efficient path to the remote network. It is also possible for a packet to exit one Internet link and then return through the other link because of the remote network's routing policy. This routing policy element has been accounted for in this design by controlling return traffic using the IGP. The active preferred path is symmetrical because of IGP costing.

Failover between the links is an automatic process because of BGP. Once a link is lost, the propagation of the routes from the core network is lost also. This path is removed from the Internet and the other path routes all traffic.

The default route is generated when a router has an active BGP connection. The route is generated locally as an aggregate and as a reject route. Because the edge routers have a complete routing table to the Internet, the reject route drops any traffic that is *not* destined to an active BGP route. This action preserves bandwidth because if the route does *not* exist in BGP, there is no reason to forward the traffic. If the EBGP peering connection is lost, then the router no longer distributes the default route into the IGP.

The edge routers contain aggregate routes for all of the public IP addresses used on the firewalls. They are sent out to the edge routes connected through ISP, ensuring that the routes are sent to the Internet through BGP, even in the event that these routes are *not* received through OSPF. It would be possible to send them out in the event they were received through OSPF only. However, it is possible that the routes would flap if they were not received from the firewalls. Flapping routes to an ISP could cause the routes to be blacklisted on the Internet. The only case where an administrator would want to send routes, if they were received through OSPF, is if the IP address ranges would be transported to a different data center or Internet connection location on the network.

Edge Routers

In this design, two Juniper Networks M Series Multiservice Edge Routers were used. They were selected for two primary reasons: interface capacity and throughput. As shown in Figure 18, each router uses six total links. Table 4 summarizes the connectivity characteristics of each link for both the A and B edge routers. For a description of the conventions used for Juniper Networks devices and links, see *Appendix B Naming Conventions.*

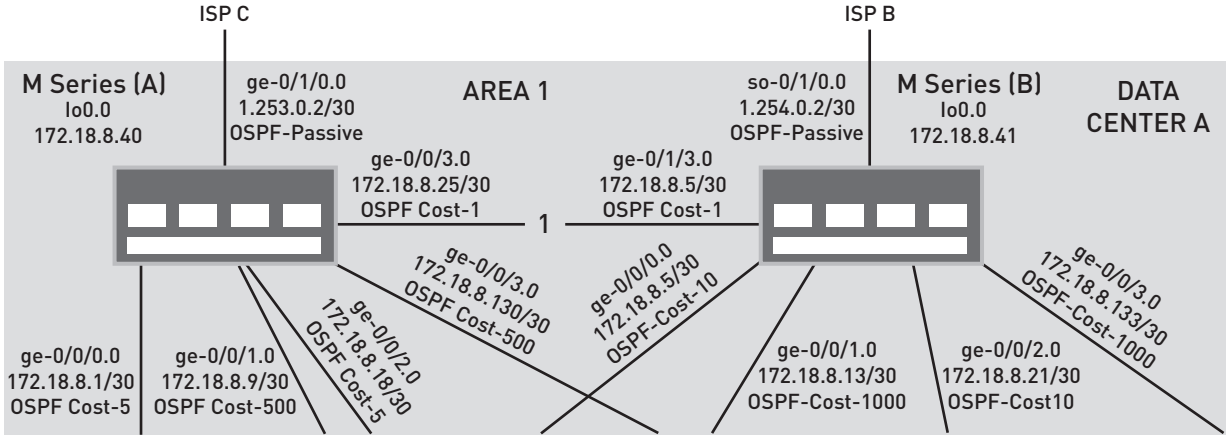

Figure 18: M Series Multiservice Edge Routers

Copyright © 2010, Juniper Networks, Inc.

Table 4: M Series Edge A and B Router Connectivity and Interfaces and Gigabit Ethernet Ratios

Function Interfaces	M Series (A) Router	M Series (B) Router
Internet provider interface	Gigabit Ethernet (ge-0/1/0.0)	SONET OC-12 (so-0/1/0.0)
Internet firewall SSG Series (A)	Gigabit Ethernet (ge-0/0/0.0)	Gigabit Ethernet (ge-0/0/0.0)
Internet firewall SSG Series (B)	Gigabit Ethernet (ge-0/0/1.0)	Gigabit Ethernet (ge-0/0/1.0)
IPsec VPN interface ISG Series (E)	Gigabit Ethernet (ge-0/0/2.0)	Gigabit Ethernet (ge-0/0/2.0)
IPsec VPN Interface ISG Series (F)	Gigabit Ethernet (ge-0/0/3.0)	Gigabit Ethernet (ge-0/0/3.0)
Router-router interface	Gigabit Ethernet (ge-0/1/3.0)	Gigabit Ethernet (ge-0/1/3.0)

Each router has a single connection to the Internet. The design uses a Gigabit Ethernet connection to connect the A router to the Internet, while the B router link uses an OC-12 connection. Connectivity from the edge devices to each of the four firewalls creates a fully meshed network. Each router employs a single connection to each of the four firewalls. Both routers use a 4-port Gigabit Ethernet PIC to create the mesh. The PIC is oversubscribed by a 4:1 ratio and should be adequate for this design, considering that a majority of the traffic is destined for the VPN firewalls. If oversubscription occurs, a second Gigabit Ethernet PIC can be added if throughput performance is insufficient.

The edge routers are linked to each other through a single Gigabit Ethernet link that provides a transit path around a less preferred or failed path. The Ethernet link also provides continuity to the Internet so both Internet connections can be used for the best possible path to the remote networks and in case a packet is asymmetrically routed over the Internet and returns over the less preferred link. Because the routers provide only a default route to all the SSG Series firewalls, the SSG Series firewalls will only send traffic to the next best router. The routers then must determine which of the two Internet links to use. If the router with the less acceptable path receives the traffic, it will send the traffic to the router that contains the preferred path.

The link connecting the router to the Internet uses a passive interface. An OSPF passive interface allows its network to be redistributed but no neighbor relationships can be established. The other five links are incorporated into OSPF and are weighted to ensure that traffic flow occurs in a predetermined, predictable pattern. Doing so simplifies troubleshooting and ensures that traffic reacts only as intended. Figure 19 shows the OSPF costing relationships assigned to the links.

Each of the firewall clusters has a primary and a secondary firewall and the links favor the primary firewall. In the event of a primary firewall failure, the secondary firewall takes over, even though the paths have higher costs. Higher costs paths are chosen because they are the only paths available. This guarantees that traffic traverses the same path in both directions. The advantage of using cross links is that traffic remains within the same firewall during a single link failure.

The OSPF connection between the routers is monitored with Bidirectional Forwarding Detection (BFD). BFD is a draft protocol that allows for subsecond detection of link failures. It integrates with the routing protocol and notifies the routing protocol when the peer is unavailable, allowing for quick convergence. In this scenario, BFD removes this path from the routing protocol within 900 ms. Without BFD, path removal typically takes a maximum of 40 seconds. Juniper Networks recommends using BFD as best practice where available.

It is critical that the routers are properly configured to ensure correct routing throughout the design. The configuration schema defines the A router as the preferred path so it is selected as the best path, and the B router is designated as the less preferred path and is used only in the case of a failure. Refer to the *Implementing a High Availability Enterprise Network* for the Data Center application note, which provides the edge router configuration parameters for both edge routers.

Firewalls—Internet and IPsec

The firewalls in this solution consist of two separate firewall deployments that are configured for separate purposes. One firewall provides secure access to the Internet. The second firewall exclusively provides a VPN termination point. Both firewalls use the concept of a mixed mode high availability deployment, which offers a combination of virtual security interfaces (VSI) and non-VSI interfaces to pass traffic.

Copyright © 2010, Juniper Networks, Inc.

The first set of firewalls, specifically the Internet firewalls, needs to provide a few specific functions. The Internet firewalls must host the NOC. The NOC is deployed off of the firewalls like a traditional DMZ to ensure that the data it collects is secured and unaltered by attackers. The Internet firewalls also must connect to the Internet and receive routing information from the edge routers. This routing information is then passed to the shared services core connected behind the Internet firewalls.

The second set of firewalls, the IPsec VPN firewalls, provide the connectivity hub for all remote sites and terminate IPsec VPNs from the Internet as well as from the private WAN. These firewalls connect to the Internet and receive routing information. This allows for connectivity to the Internet and provides access for the remote branches to connect and terminate their VPNs. The connection to the private WAN is similar to the Internet, except that the network is private. The IPsec firewalls also terminate VPN tunnels for all of the remote branches over the private WAN. To provide services into the remote branches, the IPsec VPN firewalls must connect to the shared services core.

The firewalls are designed for availability and are similar to those found in a data center or Internet perimeter. Besides the inclusion of redundant hardware, dynamic routing protocols and fully meshed links are employed to minimize the amount of failure cases that could impede Internet access.

The separation of Internet accessibility and VPN termination is a choice that each organization must make if it employs VPN services. In this solution, the separation provides a highly scalable VPN services infrastructure without dependencies on the availability of the Internet firewall. If a network flood or Internet-based attack occurred on the Internet services firewalls, it is possible to lose VPNs if the two were integrated. Therefore, the VPN devices are scaled to provide several thousand tunnels and are limited to only this function. The VPN termination firewalls terminate *not* only Internet-based VPNs but also VPNs that come across the private WAN.

Internet Firewalls

The hardware requirements for the Internet firewall are fairly basic and generally not required to provide a great deal of throughput or concurrent sessions. They are only needed to integrate with dynamic routing, secure the NOC, and secure access to the Internet. For this purpose, Juniper uses a cluster of SSG Series firewalls to implement the Internet firewalls. If more throughput and more services were going to be hosted behind the Internet firewalls, a cluster of Juniper Networks ISG Series Integrated Security Gateways would be used.

The NSRP configuration used for the Internet firewalls is a mix of active/active and active/passive. First, **VSD 0** was unset to make all of the physical interfaces have unique IP addresses. The Interface IP addresses are not synchronized among cluster members, allowing both firewalls to maintain their own OSPF neighbor relationships. If a failover occurs, the firewall does **not** need to build its own relationships, which eliminates downtime during the transition. However, this solution requires terminating interfaces for two separate purposes. The first is used for NAT (both source and destination), and the second is used as a gateway for the hosts on the NOC. To do this, the cluster must also act as a traditional active/passive NSRP cluster.

To have the firewall also provide these terminating interfaces, a **VSD: VSD 1** was created. This allows a VSD to overlay on top of the firewalls. It acts as a traditional active/passive cluster, with one firewall being the primary owner. Juniper created two VSI interfaces to accomplish this.

The Internet firewalls are designed to be resilient through two specific failure cases: interface failure and device failure. To overcome the possibility of interface failure, the firewalls use meshed links where possible. If one link fails, the second link takes over for the first. It is possible to use redundant or aggregate interfaces, but these interface types were *not* chosen for two reasons. First, as the network already allows dynamic routing protocols, it automatically corrects itself during link failure. Second, ScreenOS currently has a limitation where redundant and aggregate interfaces are unable to apply the screen feature TCP SYN cookie.

The weak point in this design is that the NOC only has a single interface per device connected to it. If a firewall loses this interface, it MUST fail over to the secondary. This is a challenge because the firewall must fail over both the routing protocol and the VSD. Because these are separate autonomous systems, the device must perform some additional actions to fail over both.

The firewall must take two monitoring steps to ensure that the routing protocol and the VSD fail over together. The VSD contains the loopback interface the firewall needs for NAT traffic. If that interface does *not* exist on the firewall that is accepting the traffic, then NAT will fail, which causes the traffic to fail ultimately.

First, because the VSI that the NOC hosts uses the failover interface, and is bound to that VSD, the firewall must monitor and verify that the NOC interface is up. The firewall must fail over the VSD if the NOC interface fails. If the interface fails, then it must force down the interfaces connected to the Internet. This action prevents the firewall

Copyright © 2010, Juniper Networks, Inc.

from accepting upstream routing information, namely the default route. If the firewall were still in the process of accepting the default route, then it would take that route and pass it to the shared services core. This would force all of the traffic through the firewall even though the required VSD is *not* active on the device.

To accomplish this, interface monitoring is enabled. Interface monitoring states that if the monitored interface changes to a specific state, then that interface must change to a different state. In this case, if the NOC or out-of-band interface fails, then both Internet-facing interfaces go down. If both Internet interfaces fail, but the NOC interface stays up, the VSD must fail over because the VSD must be on the firewall that has access to the default route being received by the edge routers. To accomplish this, **VSD 1** tracks the next-hop IP addresses on the Internet-facing interfaces. If that IP address is no longer accessible, then the OSPF neighbor relationship will be down and the link essentially fails. If both Internet interfaces fail, then the firewall will fail over **VSD 1**.

The firewalls are deployed in an NSRP cluster and cross-connect to each other using 10/100 interfaces, as shown in Figure 19. One interface is dedicated as an out-of-band management port and does not pass any traffic. It is only used for management purposes. The NOC has only one interface from each firewall and it is not full meshed, so it also uses a 10/100 interface that operates like a traditional NSRP active/passive interface. Table 5 summarizes the Internet firewall connectivity and interfaces. For a description of the conventions used for Juniper Networks devices and links, see *Appendix B Naming Conventions*.

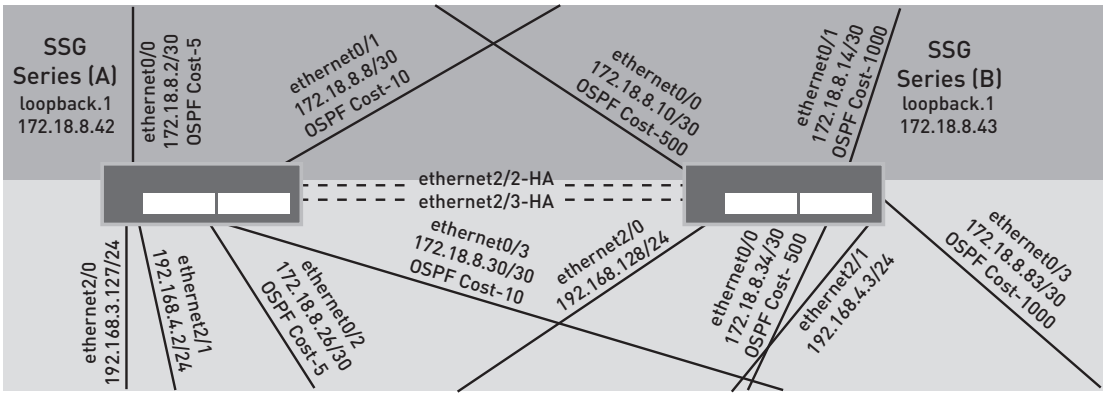

Figure 19: Internet firewalls

Each of the two firewalls has its own unique IP address on the physical interface. However, both firewalls share a VSI interface. This VSI is part of **VSD 1** and *not* the typical **VSD 0** that is used in an active/passive interface. It is a passive OSPF interface that allows only the firewall with this interface in an active state to transmit a link-state advertisement (LSA) for this network. The VSI allows all the hosts in the NOC to use this interface as a default route. If a device failure occurs, the backup device takes over using the same VSI interface.

Table 5: Internet Firewall Connectivity and Interfaces

Function	SSG Series (A)	SSG Series (B)
Edge router connectivity M Series (A)	Gigabit Ethernet ethernet0/0	Gigabit Ethernet ethernet0/1
Edge router-connectivity M Series (B)	Gigabit Ethernet ethernet0/0	Gigabit Ethernet ethernet0/1
Shared services connectivity Switch (A)	10/100 ethernet0/2	10/100 ethernet0/3
Shared services connectivity Switch (B)	10/100 ethernet2/1	10/100 ethernet0/2
Router-router HA interfaces	10/100 ethernet2/2 – HA 10/100 ethernet2/3 – HA	10/100 ethernet2/2 – HA 10/100 ethernet2/3 – HA
Router-router interface	Gigabit Ethernet (ge-0/1/3.0)	Gigabit Ethernet (ge-0/1/3.0)

Juniper deploys the other four interfaces into two separate networks as a full mesh and all use Gigabit Ethernet interfaces. Two of the interfaces connect into the Internet routers, and the other two interfaces connect into the shared services core.

Copyright © 2010, Juniper Networks, Inc.

The NSRP design uses a mix of VSI and non-VSI interfaces. NSRP is designed to make a firewall pair (or cluster) appear to operate as a single device. A master/backup protocol, NSRP, allows two devices to synchronize their configurations and operations. In the event of a failover, the backup device seamlessly picks up where the master device left off, without disrupting transit traffic or VPN services terminating on the device.

This flexibility allows the cluster to offer terminating interfaces only where needed and integrate with OSPF at the same time. It is possible to form OSPF neighbor relationships on VSI interfaces. However, if a failover occurs, excessive traffic loss will happen because the OSPF neighbor relationship must be re-established. In this solution, four of the interfaces have OSPF neighbors established, including Internet and shared services interfaces. Similar to the edge routers, these links are also weighted so that one is preferred over the other. Using the routing information learned over OSPF, the firewalls learn how to access remote networks.

Figure 20 shows the OSPF costs and relationships assigned to each of the links.

The Internet interfaces are in the untrust zone, which is contained in the untrust virtual router. This arrangement ensures that the edge routers will *not* learn the routes from the shared services core. All the interfaces and the untrust virtual router only participate in OSPF area **1**. In this case, there is no OSPF area **0** inside of the routing domain for the untrust virtual router. It is possible to run this as OSPF even if this single zone (area **0**) is not an OSPF designated area. Because area **0** does exist inside of the shared services core, this solution uses a different area to eliminate any confusion.

The untrust virtual router is configured to share some routes that it learns and export these routes to the trust virtual router. First, if the firewalls receive a default route from the edge routers, then that default route is sent to the trust virtual router to notify the firewalls where to send its traffic. Second, the loopback IP addresses of the routers and the Internet firewalls are passed to the trust virtual router. This is done for monitoring purposes because the loopback IP addresses are used for monitoring only. The loopbacks from the routers are exported via a route map that looks for the routes in OSPF. The loopbacks on the firewalls are exported via a route map as a connected route only. This approach prevents the firewalls from exporting each other's loopback IP address for monitoring.

If the router redistributes its own loopbacks from OSPF, then the router sends only the loopback from the other firewall. This results in asymmetric routing because the firewall only sees the other cluster member's route and exports it. The traffic would enter the opposite cluster member and then enter the proper cluster member on the wrong interface. As a result, monitoring fails because the firewall routes the traffic to the wrong interface.

In this deployment, the IP addresses used on all of the physical interfaces are private, non-Internet, routable IP addresses. This means that no one outside the edge routers has the ability to contact these interfaces. However, the interface **loopback 2:1**, a virtual security interface, contains a small subnet with public Internet routable IP addresses. This interface uses Mapped IPs (MIPs) as static NAT mappings for the Juniper Networks Network and Security Manager (NSM) servers and the SSL VPN NOC gateway. This interface exists in OSPF as a passive interface and is in the virtual security device number **1**.

This VSI is active on only one device at a time. In this design, VSI is active primarily on the **A** firewall. Only the firewall that has the active loopback interface sends the link state acknowledgement (containing the route to the network contained on the loopback interface). If a failover occurs, the secondary firewall activates the loopback interface and then broadcasts the LSA for that network. This makes failover quick and reduces traffic loss to only a few seconds.

The shared services core interfaces are in the data center zone, which represents everything within the data center and exists in the trust virtual router. This area participates in a separate OSPF instance and in OSPF area **0** only. This configuration separates the routing in the untrust zone from the trust zone and eliminates the private routes from being sent to the edge routers. Each of the interfaces is assigned a cost in a tiered fashion to ensure predictable traffic flow. If the untrust virtual router learns the default route and the routes from the loopback interface, they are then distributed through OSPF to all of the firewall's neighbors as Type 2 routes. Type 2 external routes are used because they take the cost of network links into account when calculating distance.

Type 1 external routes do *not* take metrics into account and they might send traffic to undesirable paths.

The shared services core interfaces are connected directly to the shared services switches. Each interface from the firewall is directly connected to a routing interface on the switches that have a 30-bit subnet on the interface. This creates a point-to-point network for the creation of an OSPF neighbor relationship. The interfaces are all configured as OSPF point-to-point links, which ensures that a designated router does not need to be elected on the interface, thus speeding up the adjacency process. The Internet firewalls learn all of the routing information from the shared services core.

Copyright © 2010, Juniper Networks, Inc.

The firewall policy that is deployed on the firewalls allows for minimal amount of access through the firewall. The firewall is configured to secure monitoring traffic as it enters and exits the NOC. It allows traffic from the Internet to access the two configured MIPs. These MIPs support the remote branch firewalls and allow them to be managed via NSM. The firewall also allows traffic to exit the data center for a minimal amount of services. MIPs are primarily for update services for the servers in the data center. This outbound traffic is NATed through a Dynamic Internet Protocol (DIP) located on the loopback VSI.

The Internet firewalls are configured with screening to reduce the threat from network-based floods. In testing the screen configuration, it was able to easily ward off a 100 pps SYN flood attack with minimal performance impact. For a listing of the firewall policy configuration parameters, see *Implementing HA at the Enterprise Data Center to Connect to a Large Number of Branch Offices*.

VPN Firewalls

The VPN firewalls need to provide many different high-end tasks for this solution. They must do the following:

- Support a minimum of 2,000 IPsec VPNs
- Redistribute at least one route from up to 1,000 branches
- Run both RIP and OSPF
- Secure traffic by providing IPS-level scanning

The ISG Series firewall is the best firewall to do this because it can easily handle all of the aforementioned tasks. Figure 20 provides a detailed view of the ISG Series VPN firewalls.

The NSRP configuration is similar to the Internet firewalls. The VPN firewalls also are a mix of active/active and active/passive. In this case, the firewalls need to have several different terminating interfaces for VPNs instead of just NAT, or as a default gateway. We use a total of seven VSI loopbacks in this design. For the VPN firewalls, **VSD 0** was unset, as the firewalls needed to maintain their own OSPF neighbor relationships. Two different VSDs were created to allow each of the firewalls to handle some of the VPN traffic, effectively load-balancing the traffic. It is possible to load-balance traffic because of dynamic routing.

The distribution of the seven VSIs provides load balancing in two different ways—across firewalls and across the two ISP. Each of the two ISPs used allocated several small eight-host networks. These networks were divided into individual IP addresses by creating a single host subnet. This approach uses 10 full IP addresses because the subnet IP and broadcast IPs can be used as individual host subnets. For **VSD-1**, which is designed to be primary on the ISG Series (E) firewall, it contains four interfaces. The second VSD, **VSD-2**, has three interfaces and ISG Series (F) is primary for them. Table 6 illustrates the significance of the IP addresses that are deployed.

Table 6: IP Address Deployment

Interface Name	IP Address	IP Owner	Use	Primary Firewall
Loopback.1:1	1.2.0.6	ISP B IP	VPN Termination	ISG Series (E)
Loopback.2:1	1.3.0.6	ISP C IP	VPN Termination	ISG Series (E)
Loopback.32	1.2.0.7	ISP B IP	VPN Termination	ISG Series (F)
Loopback.42	1.3.0.7	ISP C IP	VPN Termination	ISG Series (F)
Loopback.5:1	172.18.8.162	Private WAN	VPN Termination	ISG Series (E)
Loopback.6:2	172.18.8.164	Private WAN	VPN Termination	ISG Series (F)
Loopback.8:1	172.18.8.169	Private WAN	NAT for NSM	ISG Series (E)

Having multiple VSIs on the same firewall with different ISP IPs balances the traffic across both of the ISPs automatically. First, the remote sites simply need to be configured in a pattern to ensure that each of the loopbacks receives the same amount of remote sites bound to them. Second, each VPN firewall has the same ISP address deployment. Therefore, the second VSD on the second firewall also can have VPNs terminated to it as well, allowing traffic to be balanced across both ISPs and both firewalls.

If a failover occurs, a VSD from a failed firewall will be passed to the remaining firewall cluster member, with seamless failover similar to an active/passive cluster. The private WAN loopback interfaces are deployed for VPN termination as are the Internet VPNs. The difference is that there is one VSI per VSD for the private WAN

Copyright © 2010, Juniper Networks, Inc.

because there is only one private WAN provider. The last loopback interface, **Loopback 8:1**, is used for NAT in the same manner as the Internet firewalls. Juniper configures NAT for both NSM device servers to allow for direct communication over the private WAN.

NSRP ensures stateful failover (for active user traffic, and services between VPNs) by continuously synchronizing Real-Time Objects (RTOs) and configuration information between firewalls in a cluster. This information is sent across the HA links between firewalls.

NSRP failover monitoring is simple, and because there is no DMZ, complex monitoring scenarios do not exist. All of the interfaces participate in OSPF. The only terminating interfaces are loopbacks that exist inside the firewall, and they are not bound to a specific physical interface. Device monitoring is accomplished by ensuring that all three of the needed zones are available. If all of the interfaces bound to a specific zone fail, then the device will fail over. Because the firewall does not pass a default route from the Internet edge routers to the shared services core, the firewalls will not fail in alignment with the routing protocols. If a VSI is not available, then the remote site's tunnels fail and the backup tunnel for another VSI takes over.

The VPN firewalls use a total of eight physical interfaces. Table 7 lists the VPN firewall connectivity and interfaces. Because of the design of the ISG firewall, all but one of the interfaces is card-based. The Management Interface is an onboard interface that connects to the out-of-band management network using a 10/100 Ethernet interface. Two interfaces in slot four are dedicated as HA ports for NSRP. Slot four contains a four-port 10/00 card. A 10/100 interface has sufficient bandwidth to support state sync for NSRP. For a description of the conventions used for Juniper Networks devices and links, see *Appendix B.*

The two Gigabit Ethernet ports provide the connection to the Internet edge routers. Both of these ports are on the same card in **slot one.** Each port connects to a separate router and is weighted in OSPF. In this manner, one link is preferred over another and ensures that traffic will flow as expected. The WAN links operate the same as the Internet links on the Internet firewalls. They are in OSPF **area 1** and are a standalone OSPF area.

The VPN firewalls block traffic from the shared services core to the untrust network. The only routes that are imported from the untrust virtual router to the trust virtual router are the individual loopback interfaces used for network monitoring. These routes are distributed via OSPF to the shared services core so that the NOC systems can determine how to access the loopbacks. There are no firewall policies in place that otherwise allow traffic to leave the data center through the VPN firewalls.

Table 7: VPN Firewall Connectivity and Interfaces

Function	SSG Series (A)	SSG Series (B)
Edge router connectivity M Series (A)	Gigabit Ethernet (ethernet1/1)	Gigabit Ethernet ethernet1/2
Edge router-connectivity M Series (B)	Gigabit Ethernet ethernet1/1	Gigabit Ethernet ethernet1/2
Shared services connectivity Switch (A)	10/100 ethernet2/1	10/100 ethernet2/2
Shared services connectivity Switch (B)	10/100 ethernet2/1	10/100 ethernet2/2
WAN router connectivity J Series (A)	10/100 ethernet3/1	10/100 ethernet3/1
Router-router HA interfaces	10/100 ethernet2/2 – HA 10/100 ethernet2/3 – HA	10/100 ethernet2/2 – HA 10/100 ethernet2/3 – HA

A second card is deployed in **slot two** using two Gigabit Ethernet ports. These two ports are configured as individual Ethernet ports in a full mesh to the two switches in the shared services core. The two links have different costs to ensure that a specific link is preferred. Similar to the Internet interfaces, each of the interfaces is weighted so that one is preferred over the other. These interfaces participate in OSPF **area 0** inside of the shared services core. Figure 20 illustrates the port configuration as well as the OSPF costs relationships assigned to the links.

The private WAN is deployed from the VPN firewalls to ensure that any traffic from the private WAN terminates only to these firewalls, thereby securing the traffic. Besides using VPN, the firewalls also have integrated IPS to secure traffic inside of the VPN. This approach reduces the possibility of any threats entering the data center or from

Copyright © 2010, Juniper Networks, Inc.

access between branches. Many organizations often terminate the private WAN into the shared services core or core network. This method potentially allows attacks or allows unsecured traffic to migrate to the private WAN from the branch offices. Even though the private WAN provides a different type of connectivity than an Internet-based VPN, the traffic should still be treated the same as the Internet traffic.

Juniper Networks integrates the VPN firewalls with two routing protocols: OSPF and RIP. OSPF is the primary interior routing protocol used throughout the solution. RIP is used on the tunnels established with the various branch locations. OSPF is used to pass routes between all of the devices. However, OSPF *cannot* be used for all of the VPN tunnels because large VPN topologies do *not* accommodate the use of this protocol. First, all of the remote devices, which are small, cannot handle all of the routing information. Second, one change in the topology forces all of the sites to recalculate SPF. It is possible for this to happen several times in an hour, causing a flurry of calculations. OSPF also does not filter routing information, except at area borders.

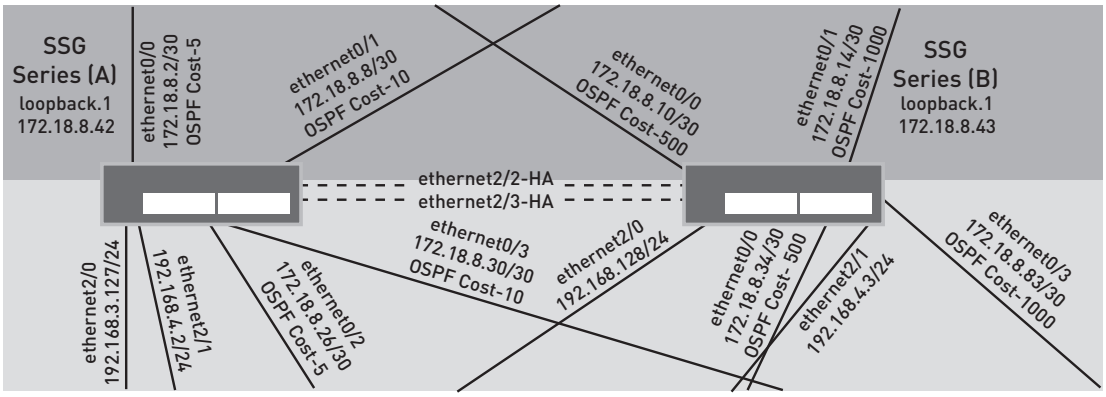

Figure 20: VPN firewalls

ScreenOS supports two other routing protocols: RIP and BGP. This solution uses RIP. While BGP is better suited for larger deployments, the protocol can be cumbersome and complex to manage. Because customers already use RIP for broadcasting routes, it was an easy decision to continue using this protocol for all of the remote branches as well. RIP is enabled on the tunnel interfaces only. All of the learned routes from the remote branches move into the trust virtual router. These routes are then distributed into OSPF as Type 2 external OSPF routes. This methodology allows the data center to determine the location for all of the remote sites. Because two firewalls are active at any time, all of the individual routes must be sent to the entire data center. See *Implementing HA at the Enterprise Data Center to Connect to a Large Number of Branch Offices* for protocol configuration settings. The RIP preference has been changed so that the learned RIP route for the remote location will be preferred over the learned OSPF route.

All of the VPNs that terminate at the firewall use the ScreenOS standard implementation for security. This means that the tunnels use 3DES encryption and SHA1 for hashing. Each of the VPN firewalls has 10 tunnel interfaces. The tunnel interfaces employ the Next Hop Tunnel Binding (NHTB) protocol to determine the remote networks that are associated with their respective VPNs and tunnel interfaces. Using the RIP feature demand circuits, the NHTB table is automatically updated to reflect the location of the remote network.

The firewall security policy is nearly identical to the policy configured at the branch office. The policies are extremely restrictive, allowing access for only the minimum required services. All of the policies on the VPN firewalls have intrusion detection and protection inspection enabled. The IPS policy is set up to stop the most critical recommended attacks. The policy also logs any minor attacks. Figure 21 shows the IPS policy used on the VPN firewalls.

	No.	Source	Destination	Terminate Ma...	Look For Attacks	Action	Notification	VLAN Tag	Install On
⚠	1	any	any	☐	[Recommended]Critical	Drop Connection	Logging	Any	DC VPN Firewalls
⚠	2	any	any	☐	[Recommended]Major	Drop Connection	Logging	Any	DC VPN Firewalls
⚠	3	any	any	☐	[Recommended]Minor	Drop Connection	Logging	Any	DC VPN Firewalls
⚠	4	any	any	☐	[Recommended]Warning	None	Logging	Any	DC VPN Firewalls

Figure 21: VPN firewall IPS policy

Copyright © 2010, Juniper Networks, Inc.

Shared Services Core

The shared services core is the section of the network where all of the network components converge. The shared services core consists of two shared services core switches, which are configured for both routing and switching. All of the interfaces connected to the firewall are routed interfaces that participate with the firewalls in OSPF **area 0**.

Several networks exist on the shared services switches. The switches offer a terminating interface for the servers by using a Hot Standby Routing Protocol (HSRP) interface. This interface is available if a switch fails. The server networks are distributed into OSPF to allow all of the devices in the network to reach the servers. If a switch fails, then the backup switch continues to notify the network to locate the servers.

High Availability Design of Firewalls

The high availability design of the firewalls incorporates two important design elements:

- A strong integration with dynamic routing - This allows the firewalls to integrate as a router where needed in the design. It gives each firewall unique IP-addressed interfaces for interacting, using the OPSF protocol.
- The use of VSI where needed - The VSI is used as a shared interface between the two firewalls, allowing a single interface to represent the cluster. Because the clusters are using a mix of VSI and non-VSI interfaces, it is called a mixed-mode cluster.

Dynamic routing determines the flow of the traffic path in this solution. The design provides a dynamically available environment. Each tier is deployed as a fully meshed solution. This provides redundant paths on each redundant device. If a link fails, a single device is **not** lost, thereby increasing the opportunity for uptime in the environment and avoiding removal of a viable path.

In case of a failure, the network requires a minimum of one additional redundant path to route around the failure. While this design offers HA, adding a second data center further enhances HA because you could lose an entire data center, yet still have network operability.

Quality of Service Design Requirements

The following are the quality-of-service (QoS) design requirements associated with this implementation:

- When more than one service provider is used, each provider may require a different Differentiated Services code point (DSCP) value for each class of service. In this case, interfaces connecting to different service providers should be assigned to different zones (*Untrust1* and *Untrust2*). This way, you can configure different policies for traffic designated to each of the different providers so different DSCP values can be used, depending on the destination zone (and therefore the destination provider).
- Virtual channels are supported only on J Series platforms. Regional offices/data centers using M Series or T Series routers *cannot* provide per-branch queuing/shaping.
- When using Juniper Networks WX Series Application Acceleration Platforms, QoS is enforced by WX OS. This effectively replaces virtual channels at the regional offices as the WAN accelerators enforce a maximum bandwidth on a per-tunnel basis.
- Traffic marking currently is not supported on AC VPN tunnels.

Copyright © 2010, Juniper Networks, Inc.

WX Design Requirements

Table 8 summarizes the WX Series design requirements. For detailed information pertaining to the WX Series Application Acceleration Platforms, refer to *WX Series/WXC Series WAN Acceleration: Implementing WAN Acceleration at the Branch Office* application note.

Table 8: WX Series Design Requirements

Requirements	Description
IPsec termination	Branch offices will encrypt all traffic sent to and received from the data centers/regional offices.
IPsec split tunneling	IPsec split tunneling splits the tunnel into two paths. One goes directly to the Internet or as clear text—the other goes over an encrypted link to the data center or the regional office.
UTM services	Traffic originated and terminated at a particular branch office must be inspected and processed by the firewall, Deep Inspection, antivirus, and content filtering modules.
Branch-to-branch communication	Branch-to-Branch communication uses IPsec tunnels (using either manual configuration or Auto-Connect VPN). However, it is assumed that most of the branch-to-branch traffic will be used for voice communications and therefore does not need to be optimized.
Traffic classification and prioritization	Traffic must be classified and prioritized before being sent out to the WAN. See Quality of Service Implementation Guide for more details.
HA	It must be possible to integrate the WX Series devices in environments where high availability (HA) is required. Refer to Implementing HA at the branch for more information.

Summary

This guide provides Juniper Networks design guidance and recommendations for creating a dynamically and highly available network extending from the data center to branch locations. It provides direction in designing and deploying an enterprise network that supports a maximum of 1,000 branch office locations and offers a dynamically available network solution. For practical application of these concepts, see *Implementing a High Availability Enterprise Network for the Data Center.*

Appendix C lists the product tables for the various Juniper Networks and Juniper partner product solutions that support the design and deployment of high-performance enterprise networks.

Appendix A Related Documents

The following table provides a consolidated list of documents centric to this design.

Document Title
Juniper Networks Enterprise Framework
Branch Office Reference Architecture
Implementing IPsec VPN for Branch Office Connectivity Using RIP
Implementing HA at the Branch Office
Implementing HA at the Enterprise Data Center to Connect to a Large Number of Branch Offices
WX Series/WXC Series WAN Acceleration: Implementing WAN Acceleration at the Branch Office

Copyright © 2010, Juniper Networks, Inc.

Appendix B Naming Conventions

M Series and J Series Interface Naming Conventions

The following provides the naming conventions used for all of the interfaces on the M Series and J Series routers.

M Series Interface Names

On the M Series Multiservice Edge Routers, when you display information about an interface, you specify the interface type, the slot in which the FPC is installed, the slot on the FPC in which the PIC is located, and the configured

port number.

In the physical part of the interface name, a hyphen (-) separates the media type from the FPC number, and a slash (/) separates the FPC, PIC, and port numbers:

```
type-fpc/pic/port
```

Logical Part of an Interface Name (M Series Routers)

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers:

```
type-fpc/pic/port.logical
```

Example: "ge-0/1/0.0" indicates a Gigabit Ethernet connection - FPC 0, PIC 1, Physical Port 0.Logical Unit 0

J Series Interface Names

On the J Series Services Routers, when you display information about an interface, you specify the interface type, the slot in which the Physical Interface Module (PIM) is installed, 0, and the configured port number.

In the physical part of the interface name, a hyphen (-) separates the media type from the PIM number, and a slash (/) separates the PIM, 0, and port numbers:

```
type-pim/0/port
```

Example: "fe-0/1/0.0" indicates a Fast Ethernet connection - PIM 0 Slot 1/0/Physical Port 0.Logical Unit 0

SSG Series and ISG Series Interfaces

The SSG Series and ISG Series devices include integrated 10/100/1000 interfaces. It does not use a modifier to indicate the distinction between 10/100 and Gigabit Ethernet connections. The SSG Series and ISG Series naming convention is simplified as shown in the following:

Example: "ethernet0/0" indicates Ethernet interface 0/port 0

Copyright © 2010, Juniper Networks, Inc.

Appendix C Products

Product Tables

Product/Technology	Type A - Basic	Type B - Optimized	Type C – Critical
Firewall/VPN with full UTM features (antivirus, anti-phishing, anti-spyware, anti-adware, anti-keylogger, anti-spam, Web filtering)	Uses integrated security functionality in the SSG Series	Uses integrated security functionality in the SSG Series	Uses integrated security functionality in the SSG Series
Routing	Uses integrated routing functionality in the SSG Series	Uses integrated routing functionality in the SSG Series	Uses J Series
Unified threat management	Integrated into SSG Series	Integrated into SSG Series	Integrated into SSG Series
Unified access control	Policy enforcement is integrated into the SSG Series firewall. Juniper Networks IC Series Unified Access Control Appliances are typically deployed at the data center or headquarters location.		
WAN application acceleration	— —	WX Series/WXC Series	
Centralized management	NSM for security Juniper Networks WX Central Management System for WAN application acceleration Junos Scope Software for enterprise routers		
Centralized reporting & control	Junos OS		

Product Recommendation Based On Branch Office Profile

	Small	Small-Medium	Medium	Medium-Large	Large
	10 Users/BB	50 Users/T1-E1	100 Users/ MLPPP-MLFR	250 Users/ MLPP-MLFR	500 Users/DS3
Type A - basic	SSG5	SSG20	SSG140	SSG320/350	SSG520 or ISG1000
Type B - optimized	2 x SSG5	2 x SSG20	2 x SSG140	2 x SSG320/350	2 x SSG520 or 2 x ISG1000
Type C - critical			2 x SSG140 and 2 x J2350	2 x SSG320/350 and 2 x J2350	2 x SSG520 or 2 x ISG1000 and 2 x J4350

Partner Products

Symantec

Juniper Networks has teamed with Symantec Corporation to leverage its market-leading anti-spam solution for Juniper's small to medium office platforms to help slow the flood of unwanted email and the potential attacks they carry. Part of a complete set of UTM features available on Juniper Networks firewall/VPN gateway, the anti-spam engine filters incoming email for known spam and phishing users to act as a first line of defense. When a known malicious email arrives, it is blocked and/or flagged so that the email server can take an appropriate action. Anti-spam is available on the SSG Series as an annually licensed feature.

Copyright © 2010, Juniper Networks, Inc.

Kaspersky

By integrating a best-in-class gateway antivirus offering from Kaspersky Lab, Juniper Networks integrated security appliances can protect Web traffic, email, and webmail from file-based viruses, worms, backdoors, trojans, and malware. Using policy-based management, inbound and outbound traffic can be scanned, thereby protecting the network from attacks originating from outside the network, as well as those that originate from inside the network. Unlike other integrated antivirus solutions that are packet or network signature-based, the Juniper and Kaspersky solution deconstructs the payload and files of all types—evaluating them for potential viruses—and then reconstructs them, sending them on their way.

The Juniper and Kaspersky solution detects and protects against over 100,000 viruses, worms, malicious backdoors, dialers, keyboard loggers, password stealers, trojans, and other malicious code. Included in the joint solution is a best-in-class detection of spyware, adware, and other malware-related programs. Unlike some solutions that use multiple non-file-based scanners to detect different types of malware, this solution is based upon one unified, comprehensive best-of-breed scanner, database, and update routine to protect against all malicious and malware-related programs. Antivirus is available on the NetScreen-HSC, NetScreen-5GT Series, and the SSG Series as an annually licensed feature.

SurfControl and Websense

All Internet content that is read, sent, or received carries inherent risks. Employee access to the Internet continues to introduce new dangers and content that can negatively impact your company in four fundamental ways:

- Security Threats: Viruses, spyware, and other malware can all enter your network through web-based email, file downloads, instant messaging, PTP applications, and other non-work-related sites.
- Legal Threats: Inappropriate content can lead to gender, minority, or religious harassment and discrimination. Illegal downloading and distribution of copyrighted or illegal material over your network has legal liability issues as well.
- Productivity Threats: The temptations of non-work-related Web destinations are endless. Just 20 minutes of recreational surfing a day can cost a company with 500 employees over $8,000 per week (at $50/hour/employee).
- Network Threats: An employee can crash your network just by logging into the wrong website. Other activities, such as recreational surfing and downloading MP3 files, can divert valuable bandwidth from critical business needs.

To regulate inappropriate Web usage, Juniper Networks has teamed with both SurfControl and Websense to provide either an integrated (on-box) or redirect (two boxes) web-filtering solution.

- Integrated Web filtering: This leverages an "in the cloud" architecture hosted by SurfControl's certified hosting partner. It allows enterprises to build Web access policies from the largest URL database (over 6 million pages) spread across more than 40 categories. From the WebUI or Network and Security Manager, an administrator can assemble firewall policies that incorporate and enforce Web access rights. Integrated Web filtering is available on the NetScreen-HSC, NetScreen-5GT Series, NetScreen-25/50, and the SSG Series as an annually licensed feature.
- Redirect solution with SurfControl or Websense: Traffic is redirected from any of the firewall/VPN appliances to a customer-hosted server running the web-filtering software where Web access grant/deny decisions are made and executed. The customer is responsible for the server, the software, and the associated management of the solution. Redirect Web filtering is supported across the entire product line.

Avaya IG550

The Avaya IG550 Integrated Gateway provides an additional choice in the Avaya line of Media Gateways. Enterprises can now consolidate the number of devices that they deploy and manage in their branch offices. By embedding Avaya Media Gateway functionality in Juniper Networks J4350 and J6350 Services Routers, Avaya and Juniper can offer enterprises a one-box telephony, routing, and security solution. This solution provides high-sustained network performance when under load, integrated voice and data security, and multilevel business continuity options. This best-in-class solution is available through Avaya direct channel and certified Avaya and Juniper resellers.

The Avaya IG550 Integrated Gateway consists of two primary components: a Telephony Gateway Module (TGM) and Telephony Interface Modules (TIMs).

Copyright © 2010, Juniper Networks, Inc.

The TGM550 module inserts into any slot in the J4350 or J6350 router and delivers a rich telephony feature set to the branch office. This feature set includes:

- Central Avaya Communication Manager (and other communications applications) access
- Call center agent support
- 6-party meet-me conferencing
- Local survivability in the event of a WAN failure
- Local music-on-hold and voice announcements
- Full encryption of voice traffic

The TGM operates as any other Avaya H.248-based gateway and includes a two analog trunk/two analog station module, modular Digital Signal Processors (DSPs), and a memory expansion slot.

There is a choice of several TIMs with analog: T1/E1/PRI and BRI options. The TIM514 analog module contains 4 trunks (FXO) and 4 stations (FXS); the TIM510 DS1 module supports T1/E1 and ISDN PRI; and the TIM521 module supports 4 ISDN BRI interfaces.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net.**

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2010 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

3500143-002-EN May 2010

Printed on recycled paper